Pay-Per-Percentage of Impressions: An Advertising Method that is Highly Robust to Fraud

Joshua Goodman
Microsoft Research
One Microsoft Way
Redmond, WA
425-705-2947
joshuago@microsoft.com

ABSTRACT

In this paper, we describe a simple method for selling advertising, pay-per-percentage of impressions, that is immune to both click fraud and impression fraud. We describe assumptions required to guarantee the immunity, which impact the design of the system. In particular, ads must be shown in a truly random way, across the percentage of impressions purchased. We describe prefix-match: a system that is similar to broad-match, but more compatible with pay-per-percentage. We show how to auction pay-per-percentage matches, including prefix matches in a revenue maximizing way. Finally, we describe variations on the technique that may make it easier to sell to advertisers.

Categories and Subject Descriptors

K.4.4 [**Electronic Commerce**]: Payment Schemes

General Terms

Economics, Security.

Keywords

Fraud Detection and Prevention, Mechanism Design, and Affiliate Model. Pay-per-percentage. Advertising. Pay-per-click. Cost-Per-Click, Pay-per-Impression.

1. INTRODUCTION

> Click fraud is "the biggest threat to the Internet economy… Something has to be done about this really, really quickly, because potentially it threatens our business model."
>
> -- George Reyes, Google's CFO

Click fraud and other kinds of fraud in online advertising markets are often considered one of the largest potential threats to the online advertising market, especially the market for online search terms. These attacks have been going on for many years [7], and some of the attacks are potentially quite sophisticated [1]. Recently, Google has listed click-fraud as a significant worry, and has admitted to regularly needing to pay refunds because of it [4]. Google has also been the attempted victim of blackmail: Michael Bradley threatened to release a program that would make click fraud easy to conduct and hard to detect. Others, have not attempted blackmail, but simply sell click-fraud enabling products [6].

One way to prevent click fraud is to revert to a pay-per-impression model. In this way, clicking on an ad does not defraud the advertiser (since the advertiser is not charged any additional amount.) However, a strict pay-per-impression model is subject to its own form of fraud, impression fraud, which we will describe later.

In this paper, we show that a simple system, pay-per-percentage of impressions, is immune to both click fraud and impression fraud. In this system, an advertiser picks a keyword, e.g. "cameras" and purchases, perhaps through bidding, a certain percentage of all impressions for that keyword. For instance, an advertiser might pay $1.00 to MSN Search. In return, the advertiser might receive 10% of all impressions for "camera" for 1 week. What does this mean? It means that for 1 week, one out of ten times that someone searches for the word "camera", they will see the ad. If someone clicks on the ad, the advertiser is not charged any extra money. It might seem that this method is subject to impression fraud: someone could create fake searches for the word "camera" causing fake impressions. Notice that if there are R real impressions over the week, and F fake impressions, that the advertiser will receive $.1 \times R$ real impressions and $.1 \times F$ fake ones. The number of real impressions the advertiser receives is not affected by the number of fake impressions.

This simple observation, that pay-per-percentage-of-impressions is essentially immune to fraud, is the basis of this paper. Despite its simplicity, there are a number of important details to make this work.

First, we must describe details that ensure that the system is actually immune to fraud. For instance, as we will describe, we must make sure that the impressions are done truly at random, rather than, for example, in rotation.

Second, we must work out details about how this method works in more complex situations, such as when we want to sell, say, a broad match for a phrase. As we will show, it is difficult to sell conventional broad match ads in a pay-per-percentage scheme. Some obvious schemes actually expose advertisers to subtle fraud. We will instead describe a variation on broad match, prefix match,

ACM E-Commerce'05, Workshop on Sponsored Search Auctions, June 5, 2005, Vancouver, BC, Canada. Copyright 2005 ACM 1-58113-000-0/00/0004…$5.00.

that lets us sell pay-per-percentage matches robustly. We will describe an auction system that combines exact pay-per-percentage matches, prefix pay-per-percentage matches, and traditional pay-per-impression or pay-per-click ads in a revenue maximizing way.

Third, we must convince advertisers to use this system. We will suggest a variation on pay-per-percentage-of-impressions that may be more palatable to advertisers.

Finally, we also describe how to apply this system to affiliate advertising, which can provide some immunity against click fraud and impression fraud.

2. TYPES OF TRADITIONAL FRAUD

It is useful to review traditional types of fraud, in order to show problems with existing market types (pay-per-click, pay-per-impression), some of which may not be obvious.

The simplest traditional market type is pay-per-impression. In this market type, an advertiser agrees to pay a certain amount per impression, for a certain keyword, perhaps up to some maximum total price. The price may be determined through an auction or through negotiation. A competitor to the advertiser may generate fake searches for a keyword, in order to defraud the advertiser. This can exhaust the advertiser's budget, or lower his return on investment below profitability. Lowering the advertiser's return on investment can either cause him to drop out, or to lower his bid, reducing the cost to the fraudster to participate in the market. It might seem that impression fraud also hurts the fraudster, whose ad may be affected by his own impression fraud, but a fraudster can arrange things so as to not be impacted. For instance, in some periods, the fraudster can advertise, and not engage in impression fraud, while in others, he does not advertise, but does engage in impression fraud. Or if there is a daily budget, he can initially not be in a market. At midnight, he engages in impression fraud until his competitor's budget is exhausted. Then he enters the market. Or the fraudster may advertise on some keywords for a product, but not others, and engage in impression fraud that way. Or if the advertiser uses broad match terms, the fraudster may choose otherwise unlikely or irrelevant phrases with the term, and engage in impression fraud on them. Or if location constraints are allowed, the fraudster may use a location constraint that prevents his ad being shown in, say, South Dakota. He then engages in impression fraud through servers in South Dakota, exhausting the competitor's budget. Finally, if an advertiser creates an ad for a search term where the fraudster has the best ranking web page (e.g. a trademarked term), the fraudster may not advertise at all on that term, but simply engage in impression fraud to remove his competitor [2].

More common than pay-per-impression is pay-per-click, which leads to click-fraud. Click-fraud is the best known kind of fraud today. Like impression fraud, click fraud can exhaust a competitor's budget or lower his return-on-investment, causing him to lower his bid. The hardest part of impression fraud is making sure that a competitor is hurt more than the fraudster: with click fraud, this is trivial (don't click on yourself.)

A common variation on the pay-per-click model adjusts the pricing or the positioning based in part on click-through-rate. Some services, like Google, also have a minimum click-through rate. This creates its own additional motivations for impression fraud. Since click-through-rate is roughly #*clicks*/#*impressions*, impression fraud of various sorts becomes potentially valuable. For instance, by engaging in impression-fraud, but not click fraud on a competitor, a fraudster can lower the competitor's click-through-rate, which may effectively lower their bid or lower their position or cause them to fall below the minimum click-through-rate. This kind of fraud has reportedly occurred [2], and can be even easier than other kinds of impression fraud: the fraudster can also cause impressions for himself, but occasionally click on his own ads, thus maintaining, or even raising his own click-through rate. In addition, the fraudster can engage in a particularly devious form of subterfuge: he can detect and demand a refund for his own fraud! Thus he can get whatever benefits accrue to an increased click-through rate, without paying the price.

3. MECHANISM DETAILS

In this section, we will derive progressively more realistic versions of the pay-per-percentage model. First, we will describe why truly random ad placement is necessary, despite the higher variance it causes. Next, we will explain the problems that broad match systems introduce, and describe two variations that solve these problems, one of which is called prefix match. We will then describe an auction system for selling both exact match and prefix match pay-per-percentage ads, and then we will describe how to combine conventional pay-per-percentage ads with pay-per-impression or pay-per-click ads. We will briefly describe how to integrate multiple ads per search, and then we will describe an unfortunate new kind of fraud, misinformation fraud, in which a fraudster tries to convince a competitor to bid the wrong amount. Finally, we will compare pay-per-percentage to pay-per-impression, explaining why pay-per-percentage is better.

3.1 Truly Random

It might seem that the best way to shows ads in a pay-per-percentage system is in rotation, e.g. if I have 50% of the ads, my ad is shown every other time. This minimizes the variance in the percentage I actually receive. However, in order to keep the pay-per-percentage-of-impressions mechanism immune to fraud, it is important that ads be shown purely at random. Otherwise, an attacker may be able to engage in some form of impression fraud. Imagine there are 100 bidders for the word "digital", each buying one percent of the impressions. Imagine that the search engine displays each ad in succession, in the same order, each time. That is, on the first search, it shows the ad for advertiser 1. On the second search, it shows the ad for advertiser 2. On the 101st search, it shows advertiser 1 again. Etc. Now, advertiser 1 might wish to defraud his competitors. Occasionally, he performs a query for the word "digital" and looks at the ad. Let's say that the ad is for advertiser *i*. Advertiser 1 immediately performs 100-*i* queries. He then waits a certain amount of time. If any legitimate user queries the word during that time, the user will receive an ad for advertisement 1. Thus, advertiser 1 can make sure that much more than 1% of the time, his ad is shown. Showing the ads at random defeats this.

An alternate technique that might appear to have lower variance is to pick a random order each time. That is, we pick a random order for the 100 ads, and show each of them as queries come in. When we have shown each of them once, we pick a new random order. Our rogue advertiser will thus not know how many ads to query: he may hit his own ad. However, consider a simple case with two advertisers, 1 and 2, each with 50% of the volume. Advertiser 1 can now defraud advertiser 2 as follows: he keeps performing queries until he sees two "2s" in a row. This can only happen if the first two was part of one rotation and the second two was part of the next rotation. The next ad must be a "1". Advertiser 1 waits for a short time, hoping a real query will arrive (displaying his ad), and then repeats his procedure. He thus receives more than half of the ads.

We thus suggest picking the advertisements to show independently at random at each time. In particular, we follow this very simple procedure. Each advertiser purchases x_i% of a keyword. All advertisements are exact match. At each time, with probability x_i /100, advertiser *i*'s advertisement is shown. Since the probabilities are, by assumption, independent, no adversary can change the expectation of the number of times that advertiser *i*'s ad will be shown.

If we do care about variance, and are willing to accept a small risk, we can consider techniques that estimate the expected total traffic per period and use a shuffling method over the total expected traffic. Such a method will be subject to attacks in which the fraudster creates several fake impressions, and then, if the victim accounts for a smaller percentage than expected of impressions, creates a large number of additional impressions, exhausting the shuffle. If the shuffles are large enough, this is probably not a practical attack.

3.2 Broad Matches

Ideally, we would allow both broad and exact matches. As a precursor to combining them, we will consider a pay-per-percentage system with broad matches alone. Broad match pay-per-percentage is substantially more complex than exact match pay-per-percentage. Consider two advertisers, one of whom has purchased 80% of the traffic for "digital" and another of whom has purchased 80% of the traffic for "camera." But if 100% of all searches containing "digital" or "camera" are for "digital camera", there is no way to meet these constraints. We might try to avoid this problem by using estimates of the relative traffic of various words and phrases. For instance, if we know that there are typically 100 searches for "digital camera" and 400 searches for "camera", we could sell 100% of the matches for broad-match "digital" to one advertiser, and 80% of matches for broad-match "camera" to another. Unfortunately, even assuming our estimates based on historical data are correct, this opens us up to fraud. Imagine a clever advertiser. Long ago, he noticed that in a given period, there were 100 (real) searches for "digital camera", and no searches at all for the word camera alone. Suspecting a competitor might bid on "camera", he begins creating 400 fake searches per period for the phrase "camera". He then purchases 100% of broad-matches for the word "digital." His competitor arrives, and purchases 80% of the broad matches for "camera." We think we can meet these constraints by giving all 100 occurrences per period of "digital camera" to the first advertiser, and all searches for "camera" alone to the second advertiser (of which there are 400 in a typical period, but all of them fake ones caused by the evil first advertiser.) The first advertiser has thus been able to defraud the second.

Because of this kind of problem, we will not sell traditional broad match advertisements. There is, however, a set of variations that are immune to fraud. In particular, if the algorithm for picking a match between search phrases and keywords is independent of the possible actions of a fraudster, then we are still immune to fraud. For instance, we could use an alphabetic method: choose the first word in a message in alphabetic order, and then choose from the broad matches for that word. Or choose the most valuable word, with the values published before bid-time, and value estimated according to some heuristic carefully chosen to be difficult to influence. One of our favorite methods, because it is easy to analyze and for advertisers to understand, is to choose the word at random. We call this random choice method "weighted-pay-per-percentage", and explain it in more detail.

With weighted-pay-per-percentage, we first choose at random which word in a phrase to target, and then select it based on the percentage of volume purchased. For instance, if an advertiser purchases a weighted 10% share of the broad matches for "camera," he will receive the following: 10% of exact match searches for "camera", 5% of the searches for two word phrases that include camera, 3.33% of the searches for three word phrases that include camera, etc. When selling such advertisements, in order to prevent fraud, we still must respect certain constraints. For instance, if we sell 70% of matches for "camera", and 80% of matches for "digital", then for a phrase like "digital camera", we can sell no more than 25% of the traffic (100% - (70%/2 + 80%/2) = 25%).

Notice that weighted-pay-per-percentage is immune to fraud, because it preserves all independence assumptions . For a given real ad, the chance that any particular advertiser is chosen depends only on the percentages that he has purchased, and is otherwise independent of the actions of any other advertiser.

Our preferred method simply always chooses the first word, or always chooses the last word. We'll call this "prefix-pay-per-percentage." In this system, rather than selling a full broad match, we sell only prefixes. For instance, "digital *" would match "digital camera" and "digital computer" but not "secure digital." This makes it easy to sell any arbitrary percentage, and to be certain that one has not oversold. In practice, for English, it is probably better to actually sell suffix broad matches. For instance, "* camera" would match "camera", "digital camera", "Nikon camera", etc. Since English usually follows an adjective noun pattern, or with noun-noun complements (e.g. "Pepperoni Pizza", "hospital patient"), the first word is typically a modifier to the second[1], the second word is typically the more relevant for advertising than the first. Some other languages like Japanese tend to have the opposite behavior. The choice of prefix match or suffix match might thus be made on a per-language basis.

For various reasons, we prefer suffix/prefix to weighted pay-per-percentage. We'll describe the system as a prefix system (which generally makes pseudo-code easier to read), although in practice it might be a suffix system, depending on the language.

3.3 An Auction System

Let's say that we would like to sell both prefix matches, and exact keyphrase matches, how would we arrange an auction? If we re-write each search for "x y z" as "x y z <END>", then we can think of an exact match keyphrase "x y z" as just being a prefix match for "x y z <END> *" We will assume such a rewriting occurs, allowing us to consider only prefix matches. For simplicity, we

[1] Who says a Ph.D. in Natural Language Processing is completely useless?

will assume a first price auction, although variations on the techniques described here could be used to create second-price (Vickrey) auctions. While there are many pros and cons of first-price auctions, one advantage is that a competitor cannot affect the price an advertiser pays for a keyword, except by actually purchasing it. This minimizes the impact competitors can have, which is part of the overall goal of the pay-per-percentage system.

We will consider bids of the following type: an advertiser bids for *a%* of a keyword, and is willing to pay price *p* per percent, up to $p\times a$ total. We will not consider budget constraints. He may make multiple bids, e.g. he is willing to pay 3 cents per percentage for the first 10 percent, 2 cents per percentage for the second 10 percent, and 1 cent per percentage for the third 10 percent.

Now, let us consider how to combine bids for *x** with bids for *xy**. (We will not consider ads of the form xyz*, etc., although the algorithm could be extended to handle these, as well.) Notice that if for any subphrase *xy** we award *a%* of the traffic, we can give at most (100-*a*)% to *x** bidders. There may be multiple bids xy_1*, xy_2*, etc. We wish to assign traffic to the subphrases in such a way that we maximize our revenue.

Here is an example: imagine that there are bids for

- $1.00 80% “digital *”
- $0.75 60% “digital equipment *”
- $0.75 70 % “digital camera *”

We can merge the bids for “digital camera *” and “digital equipment *” into a special “digital ? *” bid that competes with the “digital *” bid. This results in a set of bids like this:

- $1.00 80% “digital *”
- $1.50 60 % “digital ? *”
- $0.75 10% “digital ? *”

It’s now easy to run the bidding to assign 60% of the traffic to “digital ?” and 40% to “digital *”. Once the 60% has been assigned to “digital ? *”, we can assign the traffic for that to the “digital equipment *” and “digital camera *” bidders (who don’t compete with each other.)

In practice, we use the following algorithm, which is conceptually the same as the previous example, but aggregates bids in 1% quantiles for simplicity.

```
let virtual[x, 1..100] = 0;

for each bid for a% of xy* at a price p, in
descending order by price

       let  b  =  min(a,  100-bids[xy*])  //
don’t let bidding exceed 100%

       for i=bids[xy*] to bids[xy*]+b-1

              virtual[x, i] += p;

       bids[xy*] += b
```

We have now created a set of virtual bids for *x?** that represent the value we will get if we allocate part of the *x** traffic to bids of the form *xy**. We can run an auction using real *x** bids, plus the virtual bids, each of which is a bid for 1% of the *x** traffic at price `virtual[x, i]`. Once we have allocated traffic to these virtual bids, we can sum the traffic assigned to *x?** bids, and then run a sub-auction for each *xy** bid, up to the total amount allocated to *x?**. Another way to look at this is, for each *x**, for each percentage, we first determine how much money we could make if we allocated that percentage to bids of the form *x?**, and then allow those *x?** virtual bids to compete against real bids.

We could prove that this bidding system maximizes potential revenue. The intuition of the proof is as follows. For a given prefix *x*, and percentage of revenue *j*,

$$\sum_{i=1}^{j}\text{virtual}[x,\underline{i}]$$

represents the revenue we can get from assigning *j%* of traffic to *xy* bids. We will only and always assign traffic to *x?** bids to the extent that it exceeds the traffic we can get from *x** bids.

There is an interesting question about what period the auction should cover. A day? A week? An hour? We suggest that it should cover 1 week. In particular, it will be convenient for advertisers to be able to place a constant bid. Given that real impression volume will typically be expected to fluctuate based both on time of day and day of the week, a daily or hourly auction would require ever-changing bids. A weekly auction should typically be more stable, although we can expect some seasonal variations, and variations due to holidays etc. In Section 4, we will discuss a variation on pay-per-percentage that is more robust to these kinds of variations.

3.4 Combining Pay-Per-Impression or Pay-Per-Click

Pay-per-percentage is meant to complement, not replace, traditional methods like pay-per-impression or pay-per-click. Some advertisers may simply prefer traditional advertising types. Also, without traditional broad match keywords, it may be difficult or impossible to sell all of the available advertising, since we cannot expect advertisers to bid on every possible phrase or potentially useful prefix.

We will thus seek to combine traditional ads with pay-per-percentage. In particular, advertisers can bid either for a percentage of all advertising, or on a pay-per-impression or pay-per-click basis. We will then allocate some traffic to pay-per-percentage, and some to pay-per-click. Our goal will be to maximize our revenue. To be clear, if a pay-per-percentage advertiser wins x% of the impressions, that is x% of *all* impressions. If we choose to allocate less than 100% of all impressions to pay-per-percentage bidders, say y%, then at random, y% of the time, we display a pay-per-percentage advertisement; 100-y% of the time we display a pay-per-click or pay-per-impression ad (assuming we have a suitable ad in inventory; otherwise we show nothing.) The somewhat complex part is choosing, in advance, what percentage y to allocate to such pay-per-click ads.

We use the following procedure to perform this allocation.

```
for each keyphrase k in the search query logs
        for each traditional ad matching k
                let p = the expected revenue for the traditional ad (i.e. expected revenue per impression times number of expected impressions)
                create a virtual bid for a 100% exact match bid for k at price p.
combine these virtual bids with the pay-per-percentage bids and use the algorithm of the previous section
```

Notice that the ads created here are for exact match keyphrases, even if the traditional advertisement is a broad match phrase. Any conventional bids "won" through this mechanism are simply reserved for the traditional ads. At search time, if a traditional ad is selected, then the usual bidding process can be used in an online fashion to determine which of the traditional ads to show.

In both this subsection and the previous subsection, we did not incorporate budget constraints (except that in the previous subsection, we allowed maximum percentage bids.) Budget constraints would substantially complicate the auction mechanisms. We expect that it will be difficult or impossible to find efficient, exact algorithms with budget constraints, but that heuristic methods, such as greedy search methods over expected volumes may work well in practice.

3.5 Combining Other Constraints

If we allow other constraints, such as location constraints, things become even more complex. One way to incorporate location constraints is to pretend they are virtual keywords that always occur in a particular position, say the last position. An advertiser wanting an exact match search for "pizza" in South Dakota bids on "pizza <SOUTHDAKOTA> <END>". An advertiser wanting a location independent search bids on "pizza *". If we want to also allow prefix-match location specific matches, e.g. "pizza *" in South Dakota, things could become complex with an optimal auction method. We might need heuristic search methods to handle the allocations for such cases.

3.6 Multiple Ads per Search

Most search engines do not display a single ad per search: they show many such ads, in different positions. The different positions have different click-through rates. One way to handle this is to assign different values to the different positions, based on their average click-through rates. Advertisers would then bid for a percentage of the total of the relative value. That is, with four positions, the first position might be worth 4 points, the second worth 3 points, the third worth 2 points, the fourth worth 1 point. If twelve ads are shown (120 points total), and an advertiser has purchased 10% of all points, the ad might be shown 12 times in the fourth position, or 6 times in the third position, or 3 times in the third position, or some combination. Advertisers are now purchasing a percentage of all points for a given keyword, rather than a percentage of all impressions for the keyword.

3.7 Misinformation Fraud

While pay-per-percentage is immune to fraud in all of the usual senses, there is actually a new kind of fraud that is possible, what we will call misinformation fraud. In misinformation fraud, a fraudster attempts to deceive a competitor about the value of a particular keyword. For instance, if there is a keyword for "camera", a fraudster might generate millions of fake searches, knowing that there were only thousands of real ones. The victim might believe that 10% of impressions for this keyword is worth a large sum, when in reality, it is worth a small sum. If the fraudster knows the true values of different phrases, while the victim does not, the fraudster will be at a large advantage.

To combat this kind of fraud, it is important that search engines be able to robustly estimate the relative real traffic of different keywords. This might be done by looking at historical logs (some periods of which may be thought to not contain undetected fraudulent impressions for a particular keyword), or by sampling from a small set of trusted users, or by using the subset of traffic from certain IP addresses thought not likely to participate in fraud, such as proxy servers from large legitimate companies.

3.8 Why Pay-Per-Percentage is Better Than Pay-Per-Impression

It may not be obvious that pay-per-percentage is better than pay-per-impression. After all, in both cases, advertisers are essentially paying for impressions. In addition, for an advertiser to figure out how much to bid, they will need an estimate, from somewhere, of the expected number of real impressions. In the next section we will describe a method that is even more dependent on an accurate estimate of the number of real impressions. Since we need this estimate anyway, why not just use pay-per-impression, and let the search engine estimate which impressions are real and which are fake?

As we will explain in detail below, there are several advantages of pay-per-percentage over pay-per-impression. First, it is easier to determine the percentage of real impressions than which impressions are real. Second, there are fewer data sparsity issues in determining the real volume for a keyword than in determining the volume for a keyword for a specific advertiser. Third, pay-per-percentage puts control in the hands of the advertiser: they can choose from multiple estimates from different sources. Fourth, an advertiser who has found a profitable keyword need not worry that someone can use fraud to disrupt him.

The first reason pay-per-percentage has advantages is that the problem of determining how many impressions over an entire keyword are real is much easier than the problem of determining whether any given impression is real. For instance, we might expect that if we receive impressions from 20 distinct IP addresses for a given keyword, there should be about 100 real impressions. If we see 200 impressions, with 20 distinct IP addresses, we might estimate that only half of the impressions are real, without worrying about exactly which ones.

In addition, we can use historical data to estimate the number of real impressions. For instance, a fraudster might begin engaging in impression fraud only after a particular competitor enters the market. We may be able to use data from before we suspect fraud, e.g. because of complaints, to estimate the real data.

We might try to use similar information to estimate the overall percentage of real traffic for a given keyword, and then in a pay-per-impression system, pro-rate for all advertisers, but through techniques like entering and leaving the market, making location specific bids, etc. the fraudster can avoid impression fraud on himself, while inflicting it on others. Pro-rating for everyone does not change his relative advantage.

The third reason that pay-per-percentage is better is that it puts the control in the hands of the advertiser. A search engine can provide a variety of information to advertisers, such as historical traffic information, minimum frequency, maximum frequency, their best estimate in previous periods after proprietary fraud detection heuristics, etc. The advertiser may also contract with third parties for additional information. The advertiser can use their own domain knowledge to make a keyword specific decision. For instance, they may know that a particular historical increase was due to a news event, or they may know something about the seasonality or popularity of particular terms. In the end, it is the advertiser who directly decides how much the traffic, in whatever combination of real and fake, is worth, rather than creating a bid, and hoping that the search engine provider will correctly detect fraud.

For a similar reason, pay-per-percentage may also have advantages for search engines. Some advertisers have filed lawsuits against Google, Yahoo and Ask Jeeves, alleging that they knowingly charged for fraudulent clicks [3]. If a pay-per-percentage of impressions system gave only or primarily objective evaluations of volume (e.g. historical data about numbers of searches; number of searches from non-proxy servers, etc.) they might be able to absolve themselves of any apparent responsibility for detecting fraud.

The fourth, and perhaps biggest advantage of pay-per-percentage is that an advertiser need not worry that a competitor can come in and, through fraud, disrupt a successful situation. If I am an advertiser, and have been happily making money by purchasing 10% of the impressions for a given keyword, no competitor can affect me through fraud. I can simply leave my bid intact and continue. I still have to worry about competing offers and products, or a higher bid, but at least fraud is no longer a problem.

4. A MODIFIED SYSTEM THAT ADVERTISERS MAY PREFER

Pay-per-percentage is not nearly as radical as it may seem. For instance, if you buy a 30 second ad on a 30 minute television show, you are essentially buying $1/60^{th}$ of the viewers' attention. If the show, say a sporting event, happens to attract many more or many fewer viewers than you expected, the price is typically unchanged. Similarly, with print media, if you buy a full page ad in a 100 page magazine, you have essentially purchased 1% of the magazine. If the magazine happens to have an especially good cover that sells more newsstand copies the price you pay is not typically changed.

On the other hand, online advertisers are not used to this model, and may not be comfortable with a system that sells them a percentage of traffic. What if they don't get as many impressions as they expected? What if they buy 10% of the impressions for a keyword on Google, only to have a major new release of MSN Search cause Google's traffic to drop precipitously? What if they purchase the keyword "Britney Spears" for a week, only to have her popularity drop?

We thus suggest a minor variation on our scheme that protects advertisers from most of these problems. We call this method "pay-per-impression-max." Imagine for a given keyword we expect 100 real impressions. Instead of selling percentages of impressions, we will sell actual impressions, but with a twist. If we get more searches for the keyword than we expected, we will give them to advertisers for free, in proportion to the number of impressions the advertiser has bought. With both pay-per-percentage, and pay-per-impression-max, if an advertiser buys 10 impressions for 10 cents, and we receive 100 searches, we will show his ad 10 times, and charge him 10 cents. With both systems, if he buys 10 impressions for 10 cents, and we receive 200 searches, we will show his ad 20 times, but still only charge him 10 cents. But with pay-per-impression-max, if we receive 50 searches, we will show his ad 5 times, and charge him only 5 cents.

Pay-per-percentage max protects the advertiser against lower than expected traffic. On the other hand, in one presumably rare case, it can allow a form of fraud. If both of the following happen, an advertiser will pay for more advertising than he intended. First, the seller's estimate of the real traffic volume must be too large. Second, there must be undetected impression fraud. Even in this case, however, damage to advertisers is limited. If the seller overestimated the volume of legitimate traffic by 10%, then damage to the advertiser is at most 10%, and the damage will achieve this maximum only if at least 10% of the traffic is due to undetected fraud. Similarly, if 10% of the impressions were fraudulent, then damage to the advertiser is at most 10%, and it will achieve this maximum only if the seller overestimated real volume by at least 10%.

Overall, this system, pay-per-percentage max, is more familiar to sellers, who will see it basically as a pay-per-impression scheme, while robust to almost all fraud, with the exceptions as just noted.

There is an additional advantage to pay-per-impression-max: pricing stability. As we discussed in Section 3.3, there may be seasonal variations, or variations due to holidays, etc. in traffic. Or there may be long term trends, such as a rise in MSN Search's marketshare, or a decline in popularity for Britney Spears. With pay-per-percentage, I might bid 10 cents for 10% of Google's traffic for the keyword "Britney Spears", which initially has an estimated 100 impressions per period. Now, as MSN Search improves, and Google's marketshare erodes, or as interest shifts to Avril Lavigne from Britney Spears, the expected real traffic in this keyword might be reduced to 50 impressions per period. If I do not manually readjust my bidding, then I will end up paying twice what I should. On the other hand, with-pay-per-impression-max, the pricing is in terms of expected real impressions. Assuming the search engine adjusts the expected number of real impressions as traffic in the keyword drops, the pricing remains correct: an advertiser need not adjust their bidding as often.

On the other hand, pay-per-impression-max opens up some susceptibility to fraud: if a fraudster can cause a search engine to overestimate the volume of real traffic, they can cause an advertiser to overpay. Pay-per-percentage is less susceptible to this problem. Many of the arguments in Section 3.8 about the advantages of pay-per-percentage still apply to pay-per-impression-max.

Given that pay-per-impression-max depends on somewhat accurate estimates of the real advertising volume, and loses some of the robustness to fraud, it may or may not be a better method than pay-per-percentage. In practice the choice may depend on advertiser sentiment, and the search engine's ability both to accurately detect fraud, and to convince advertisers that they have done so.

In general, advertisers like pay-per-click, especially when they first start advertising with a particular keyword or a particular search provider. Assuming clicks are not fraudulent, pay-per-click has much less risk: they can estimate their per click conversion rate much more accurately than their per-impression-conversion rate, or, worse, their per-percentage-of-impressions conversion rate. We thus see pay-per-percentage as an alternative that should be integrated with a more traditional market. Advertisers would be given the choice to use pay-per-percentage (or pay-per-impression-max), rather than being required to use it. Initially, advertisers might use pay-per-click. The subset of advertisers who felt they were victims or potential victims of undetected fraud would be especially attracted to the alternative of pay-per-percentage or pay-per-impression-max.

One of the larger problems with pay-per-percentage is that it is hard for an advertiser to decide how much to bid. But for an advertiser who has been using a traditional pay-per-impression or pay-per-click model, we can help them with their bidding. In particular, if through pay-per-click or pay-per-impression, they have previously been receiving 10% of all impressions, we can suggest that to achieve their historical rate, they should attempt to purchase 10% of all impressions. In addition, by analyzing historical auctions, we can estimate a bid price that should let them achieve that percentage.

5. AFFILIATE ADVERTISING

Affiliate advertising refers to programs where a search engine or other provider brokers ads to third party sites. For instance, when an advertiser agrees to show ads on Google search, Google may also display those ads on third party sites. If someone clicks on the ad on a third party site, Google gives about 80% of the money to the third party [2], while keeping the rest for themselves. Affiliate advertising programs account for a huge portion of the total ad market. For Google, the moderate percentage that Google keeps from affiliates advertising accounts for about as much revenue as Google earns from Search advertising. The total revenue that passes through Google for affiliate ads, most of which is passed to the affiliate sites, dwarfs the revenue from search ads.

Unfortunately, affiliate ads also create an additional motivation for fraud: affiliate fraud. In normal click or impression fraud, the motivation is to harm a rival: to either lower their return-on-investment, or to exhaust their daily budget. With affiliate fraud, the affiliate site has a motivation to generate impressions or clicks. For each one, they receive most of the revenue. This can be even more of a problem than fraud on a search site. For instance, I have an ad for "spam conference" using Google Ads. While there are one or two competing spam conferences, neither of them advertises with Google Ads, and even if they did, computer science academics typically don't engage in fraud against competitors. Normally, I don't have to worry about fraud. However, if an affiliate site also shows ads for my conference, I now have reason to worry, when before I was fairly safe. In addition, since the traffic to my site might be generated by many different small affiliate sites, I may not be able to detect whether any individual site is engaging in fraud: there is not enough data. I simply have to trust Google to detect fraud on my behalf.

We suggest that instead of selling pay-per-impression or pay-per-click ads, an affiliate might sell pay-per-percentage-of-impression ads. Since the site does not make any incremental income from a click or impression, they cannot directly defraud advertisers. However, there is an important question of how to price a percentage-of-impressions on a particular site: 10% of impressions on my personal home page is worth much less than, say, 10% of impressions on almost any other page in the world. Various third party rating sites, like Nielsen-Net-Ratings, could be enlisted to provide volume estimates. While such rating companies are not immune to fraud, this at least creates a market in fraud detection, encouraging competition, and putting control in the advertisers' hands, without penalizing affiliates.

Given the need to estimate real traffic, pay-per-percentage methods may not be appropriate for, say, page-specific advertising, as opposed to domain or site-specific advertising. For a single page, there may not be enough volume to accurately estimate the traffic. This is even more true for transient content, such as content targeted ads on a news page or a blog.

Pay-per-percentage or pay-per-impression-max can also be the method for contracting with a site. Affiliates in pay-per-click programs often complain that the way that ad-resellers, like Google AdSense, pay them is unclear. An affiliate could receive an offer in advance for 10% of their traffic for 10 cents (pay-per-percentage) or for 1 cent per impression, up to 10 cents maximum (pay-per-impression-max). The reseller can then engage in content targeting, pay-per-click advertising, pay-per-percentage, etc. The affiliate will clearly understand their payments, rather than worrying about whether, say, the reseller's fraud detection was overly aggressive. The affiliate can sell their traffic to the highest bidder, rather than having to guess which reseller may generate the most revenue.

Fraud detection is inherently problematic in an affiliate advertising market. If the reseller is overly conservative, advertisers will complain. If the reseller is overly aggressive, affiliate sites will complain. A fraud immune system, like pay-per-percentage, may be the only way to satisfy both simultaneously.

6. CONCLUSION

Fraud is one of the biggest potential problems for online advertising. Advertisers who believe they may be victims of fraud, especially fraud undetected by the search engine, will certainly be interested in a system that is immune to most kinds of fraud. Search engines selling ads that are less subject to fraud can spend fewer resources detecting it, and fewer resources dealing with advertiser complaints.

Affiliate fraud is particularly problematic. A reseller that is overly aggressive about detecting fraud risks alienating his affiliates. A reseller who is not aggressive enough risks alienating the advertisers. The pay-per-percentage model is not appropriate for all kinds of affiliate advertising; in particular, it is most appropriate for high volume sites. However, when it is appropriate, it substantially simplifies relationships.

In this paper, we have described many details of a pay-per-percentage market. We do not picture pay-per-percentage being used in isolation, but rather in combination with traditional pay-per-impression or pay-per-click markets. We have not described all of the details of such a system, many of which depend on the details of the corresponding traditional market. For instance, budget constraints in the pay-per-click market must be respected in the combined market. We have shown that traditional products, like broad match, are difficult to implement in a pay-per-percentage system in a fraud-robust way, but by changing to a prefix-match system, we can solve most of these problems.

Whether pay-per-percentage markets are ever used in practice depends on a number of issues, such as how accurately fraud can be detected in traditional markets, and whether pay-per-percentage can be integrated in an existing system without too much sacrifice. Given the size of the fraud problem, and the robustness of pay-per-percentage to fraud, it seems like a very reasonable complement to traditional fraud detection.

7. ACKNOWLEDGMENTS

Thanks to Chris Meek, David Maxwell Chickering, Jennifer Chayes, Christian Borgs, Uriel Feige and the anonymous reviewers for helpful discussions and comments on earlier drafts.

8. REFERENCES

[1] Anupam, V., Mayer, A., Nissim, K., Pinkas, B., and M. K. Reiter. On the Security of Pay-Per-Click and Other Web Advertising Schemes. *Proc. of the 8th International WWW Conf.* 1999.

[2] Delaney, K. J. Traffic Jam: In `Click Fraud,' Web Outfits Have A Costly Problem, *The Wall Street Journal*, page 1, April 6, 2005.

[3] Delaney, K. J. Internet Firms Face Legal Test on Advertising Fees, *The Wall Street Journal*, page B8, April 5, 2005.

[4] Google, SEC Filing 10K, March 30, 2005

[5] Livingston, B. New Attacks and Defenses In Click-Fraud War, *Earthweb.com.* September 21, 2004.

[6] Lote Group, ClickingAgent.com

[7] Ward, J., April 19, 1999. Bogus Clickers Are Cheating Ad Networks. The Industry Standard Magazine