Analysis of Malicious Affiliate Network Activity as a Test Case for an Investigatory Framework

Mathew Miehling[1], William Buchanan[1], John Old[1], Alan Batey[2] and Arshad Rahman[3]

[1]Centre for Distributed Computing, Networks and Security, Napier University, Edinburgh

[2]Detective Sergeant, Computer Crime Unit, Northumbria Police

[3]Financial Crimes, Financial Services Authority, London

m.miehling@napier.ac.uk

w.buchanan@napier.ac.uk

j.old@napier.ac.uk

alan.batey1@talktalk.net

Arshad.Rahman@fsa.gov.uk

Abstract: Currently there is a great deal of literature surrounding methods that can be used to detect click-fraud, but there is very little published work on actual cases of click-through fraud. The aim of this paper is to present the details of a real-life fraud, in order that lessons may be learnt to overcome this type of fraud in the future. The paper outlines a fraud that is suspected to have included both PPC and PPS from fraudulent affiliates.

This paper describes a methodology for the investigation process of affiliate network scams, including the anonymisation of personal and location details, while providing an analysis of an actual crime. In total, the case examined resulted in an estimated loss of around £200,000 with a further estimated loss of over £200,000 if further transactions had not been cancelled.

The methods used within the scam are outlined using anonymised data, and presented to highlight the malicious activity. This included both pay-per-click and pay-per-sale scams most likely using stolen identity information. It concludes with the methods that may be helpful in possibly identifying malicious activity with affiliate networks and how a framework can be setup to investigate these crimes.

The current work involves developing an investigatory framework focused on the early detection of electronic fraud, and the work done for this paper will be used as a test case on affiliate fraud data. The future aim of the research is to completely automate the investigatory framework that will allow incident data to be processed so that the context of a crime is not lost, but that it anonymises and protects the identity of those involved.

Keywords: Affiliate advertising, Click through, Fraud, Financial Services, E-Crime

1 Introduction

With the recent proliferation of online retail, merchants are competing amongst themselves to drive customers toward their sites rather than the sites of their competitors. Many merchants are turning toward Internet-based advertising using affiliate networks. These programmes include pay-per-click (PPC), pay-per-lead (PPL) and pay-per-sale (PPS) (AffStat 2009).

Unfortunately, these programmes are often susceptible to click-through fraud. The detection and damage mitigation of click-through fraud is well documented (Bloch & Eroshenko 2004, Metwally et al. 2005, Ntoulas et al. 2006, Zhang & Guan 2008), but we feel that highlighting affiliates that are

more likely to commit fraudulent acts beforehand could significantly lessen the occurrence of these crimes. The aim of this paper is to present the details of a real-life case of affiliate fraud in order to learn methods that may help to overcome this type of crime in the future. It outlines a case in which we suspect both PPC and PPS schemes were abused by fraudulent affiliates.

With an Affiliate Network, valid Affiliates create Web pages with content-related to the products, and use this to attract new customers to the merchants. In an affiliate scam scenario, a malicious affiliate sets up a site which is then used to generate fake click-through traffic, either using a fake ID to make a purchase and qualify for PPS commission or to generate pure PPC commission. With the scale on Internet-based commerce, the possible scope of this type of fraud is large, and thus must be detected at an early stage, to stop the whole of the infrastructure for e-Commerce being compromised and causing users to distrust online retail. This paper outlines a real-life case that resulted in an estimated loss in the region of £200,000, with a further estimated loss of over £200,000 in cancelled business transactions.

There are many difficulties in bringing fraudulent affiliates to justice, and so we believe that it would be preferable to identify possible fraudsters before they have had a chance to commit fraud. A system of assigning a risk value to each affiliate in a programme could possibly eliminate a large amount of merchant loss due to fraud by identifying which affiliates are more likely to commit fraudulent acts.

2 Background

In a single-tiered PPC affiliate programme the merchant pays an affiliate (the owner of the website hosting advertisements) a commission every time a customer clicks a link to the merchant's store. However, in a multi-tiered PPC programme the merchant may pay a popular search engine £0.10 every time someone clicks on their sponsored link and then that search engine may pay an affiliate £0.05 every time someone clicks that advertising link in their software or on their website. This type of multi-tiered PPC programme is known as paid-for search and the Internet Advertising Bureau (IAB) have reported that these programmes grew by 6.8% from the first half of 2008 into 2009 to £1.06 Billion, which accounts for 60% of the total online advertising expenditure (Holton 2009). However, PPC advertising is prone to abuse and does not guarantee that the merchant will make any money from a click on an advertisement, so the commission is often much lower than that earned in a PPL or PPS programme.

A PPS programme often requires more involvement from the affiliate, and in order to earn a commission the affiliate must get the customer to (Edelman 2008b):

1. Browse to the affiliate's site.
2. Click a link to the merchant's site.
3. Make a purchase on the merchant's site within a pre-determined amount of time.

Once these three events occur, the affiliate can make a decent commission on the sale as the merchant is guaranteed to make money from the transaction.

There also exist third party companies that have created their own networks in order to become intermediaries between merchant and affiliates. These affiliate networks facilitate relationships between merchants and affiliates, enabling merchants to quickly and easily outsource the finding and management of a large number of affiliates. These affiliate networks earn money by charging the merchant a percentage of the paid commissions for the use of the network's service. For example, in an affiliate network that charges a 10% fee to merchants, if a merchant were to sell £10,000 worth of goods with a 5% commission to the affiliates that brought the customers, the merchant would be paying £500 in commission fees. On top of that, the merchant would then have to pay 10% of that (£50) to the network. One downside of using an affiliate network is the lack of direct interaction between the merchant and their affiliates. With this lack of interaction, the merchants must be extra vigilant in the detection and reporting or suspected fraudulent activities so that the networks can investigate further and take appropriate actions against the affiliate.

Figure 1 provides an overview of a typical affiliate programme scam, where valid Affiliates (AffiliateB and AffiliateC) create web pages with content related to the products in order to attract new customers to the merchants. AffiliateA, though, has set up a site whose sole purpose is to facilitate affiliate fraud either by using fake or stolen identity information to generate sales commission or to generate pure click-though commission. The money paid to AffiliateA consists of a mixture between PPC and

PPS commission.

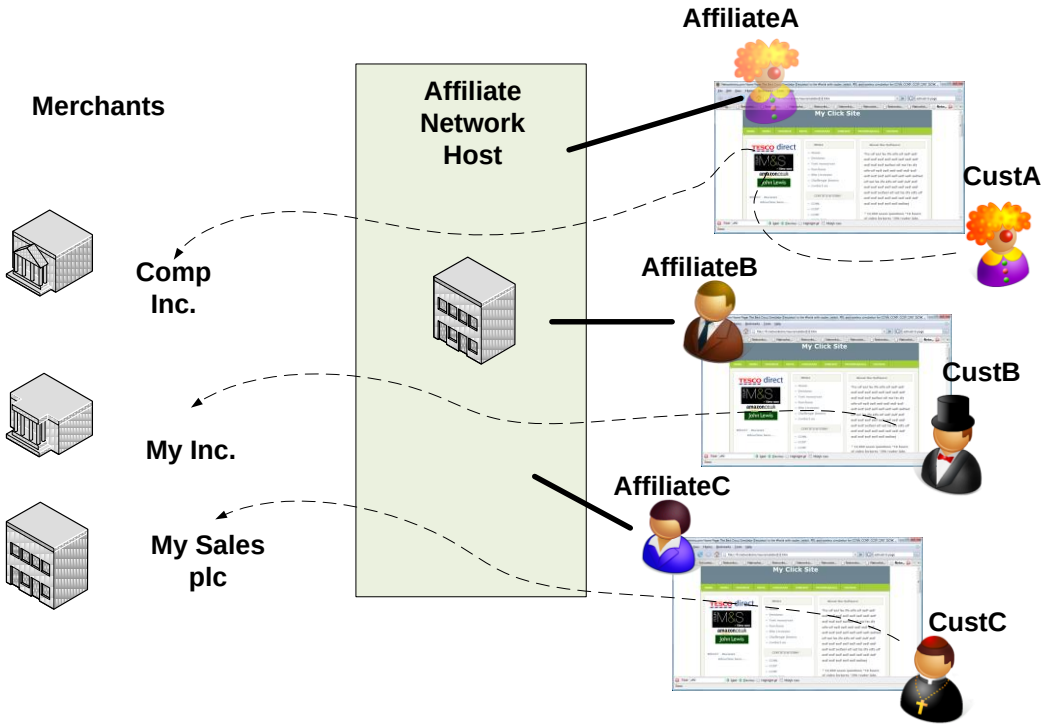

Figure 1 – Example of an Affiliate Programme Scam

3 Literature Review

Affiliate programmes offer relatively cheap and often widespread advertisement to companies, but they are also prone to abuse by malicious entities. This fraud exists in two basic forms: script-based and subterfuge. **Script-based** involves a person or script that clicks on an affiliate's link repeatedly, by request of the affiliate, with no intention of making a purchase at the merchant's site. This method is crude and easily detectable by both the affiliate network and the merchant. A simple examination of the logs would show a duplicate IP address registering clicks repeatedly from the same referrer. The address will often be from another country in which the merchant's service may not even be offered (Grow et al. 2006). **Subterfuge** employs more subtle tactics and is therefore slightly more difficult to detect. In this method, an affiliate's page sends a click to a merchant's site without the user actually clicking on anything. These fake clicks earn the affiliate money, just as a real click would, but the merchant's site is often never shown to the customer. This method is often against the Terms of Service (ToS) of many affiliate networks (Google 2009) and yet these malicious affiliates manage to slip past undetected (Tuzhilin 2006). This is because, without access to traffic logs, the actual web traffic must be monitored using a packet sniffer such as Wireshark or tcpdump to detect this method of click fraud. For a larger affiliate programme or those without staff that possess the technical knowledge to perform such tasks, this is not a feasible task. However, automated software does exist that can scan web sites and find many of the methods used to send these fake clicks on behalf of the user (Edelman 2007a). Edelman's software scours the web looking for sites that contain nefarious practices such as cookie stuffing.

Certain software vendors have created pieces of malware that exist solely to skip the first event required for a PPS affiliate (customer browses to affiliate's site) to receive commission from a sale. The malware monitors a user's web surfing sessions and when the software detects a web request for certain sites it then redirects the user's browser to the affiliate's version of the address. The necessary cookie files are then created on the customer's computer as if the affiliate had legitimately directed them to the merchant's site. If the customer then completes a commissionable action, the affiliate is paid even though they have not earned the commission (Edelman 2007b, Edelman 2008a). In order to thwart anti-virus companies, other methods of creating an affiliate's cookie on a customer's computer without first having to infect the customer with malware have been developed. Scripting languages can be employed by affiliates to simulate a click from a customer as soon as the

customer visits the affiliate's site. If that customer were to then visit any of the merchants that have had false clicks generated for them, the affiliate would wrongly receive credit for the sale and earn commission.

A slightly more effective attack eliminates both the first and second steps of the required chain of events from Section 2. An affiliate can use an invisible frame to force a false click when an ad banner is displayed in a customer's browser. This banner can then be loaded onto third party sites for a small fee incurred by the affiliate. Any customer loading the banner and then subsequently making a purchase is likely to earn the fraudulent affiliate commission. An example of how this banner works is shown in Figure 2.

The invisible frames (*italics and highlighted*) load affiliate links in new windows. Because the frames have a width and height of 0 (underlined), these windows are invisible to the user. The code to call the links is cleverly hidden in what appears to be an image file (*underlined and italics)*, but the file actually contains html calling the affiliate links. Since the image file actually contains code, the affiliate also needed to include the appropriate code to actually display an image so that the banner does not appear as a broken image.

Almost identical to the above attack, the affiliate could alternatively create an account on a forum related to the products being sold and put the invisible frame banner as their signature. Users of the forum are more likely to be in the market for the products and are likely to visit popular websites in the near future. If the customer has loaded a topic that the malicious affiliate has replied to, the cookies are placed on their system and when they make their purchase, the affiliate receives the commission.

These last two attacks are the worst possible scenario for the merchant, as they receive no advertising whatsoever because the customer is never actually told about the merchant site. Any sales that occur after this cookie stuffing would have occurred anyway but the merchant is still paying the affiliate a commission on the customer's sales.

Even more complex than cookie stuffing, some criminals have taken to using false or stolen credit card data. This data is used to make purchases in order to earn the high commission of a PPS scheme for the affiliate without costing the fraudsters any money. The fraudulent behaviour is often detected before an actual item is shipped, but it is often very difficult to bring legal actions against the fraudster due to the international nature of affiliate marketing (Zango 2005).

```
GET /iframe3? ...
...
Host: ad.yieldmanager.com
...
HTTP/1.1 200 OK
Date: Mon, 29 Sep 2008 05:36:02 GMT
...
<html><body style="margin-left: 0%; margin-right: 0%; margin-top: 0%; margin-bottom: 0%"><script
      type="text/javascript">if (window.rm_crex_data) {rm_crex_data.push(1184615);}</script>
<iframe src="http://allebrands.com/allebrands.jpg" width="468"
height="60" scrolling="no" border="0" marginwidth="0"
style="border:none;" frameborder="0"></iframe></body></html>
GET /allebrands.jpg HTTP/1.1
...
Host: allebrands.com
...
HTTP/1.1 200 OK
...
<a href='http://allebrands.com' target='new'><img src='images/allebrands.JPG' border=0></a>
<iframe src ='http://click.linksynergy.com/fs-
      bin/click?id=Ov83T/v4Fsg&offerid=144797.10000067&type=3&subid=0' width ='0'height = '0' border='0'>
<iframe src
      ='http://www.microsoftaffiliates.net/t.aspx?kbid=9066&p=http%3a%2f%2fcontent.microsoftaffiliates.net%2f
      WLToolbar.aspx%2f&m=27&cid=8' width ='0'height = '0'boder='0'>
<iframe src ='http://send.onenetworkdirect.net/z/41/CD98773' width ='0'height = '0' border='0'>
```

Figure 2 – Example of a malicious banner forcing clicks when viewing it (Edelman 2008b)

4 Methodology

The data used for this paper was provided by the Northumbria Police and includes a 2008 incident report along with a list of suspected fraudulent transactions that were cleared and a list of suspected fraudulent transactions that were stopped. The data also included a customer database containing the information related to suspected malicious affiliates. The data shows that this case involves a major incident of credit card fraud in the UK that resulted in a loss of £201,343 with a further £215,413 in cancelled payments upon discovery of the deceit.

Although the original data we received is not strictly considered live evidence, we felt it best to adhere to the Association of Chief Police Officers (ACPO) standard for digital evidence handling as laid out in (7safe & ACPO 2007). The handlers of original data should be careful that it is not changed in any way, and a copy of the data should be used for any manipulation needed. In our case, the original data is stored in an encrypted format (using TrueCrypt), and preserved in its captured format. The copy of the data used for analysis has been anonymised to ensure that personally identifying data is not revealed.

In order to remove all personally identifying information, the data needed to be anonymised. In order to adhere to the k-anonymity standard (Ciriani et al. 2007, Zhong et al. 2005), we took a blanket anonymisation approach and substituted all of the personally identifiable data with fictitious values as described in (Edgar 2004). A large list of common first and surnames was used to give each user a fictitious name, and all physical location data was changed to that of locations in movies and television shows. Table 1 shows a sample of the anonymous data used for our analysis.

We discounted sanitisation methods that simply remove identifying data, such as in (Tveit et al. 2004, Venkatesan et al. 2008), because that would make our analysis much more difficult. Working with complete, albeit fake, identities enabled us to more easily see the connections between the affiliate accounts involved.

A simple lookup table was created in order to map the anonymised records back to the original data. This is kept along with the original data, and is only to be used when mapping back to the suspect is required at the conclusion of an investigation.

The analysis that we performed was done in an abductive process. We looked at an incident report that was filed against the user identified as Stan Smith and manually followed the links from that account to several others finding more and more links between accounts along the way.

5 Case Study

According to the incident report given to us by the police, the affiliate network, AffiliateNow, received a complaint from Merchant2 about suspected fraudulent behaviour. The user involved, Stan Smith of 416 Cherry Street in Gotham City, had been sending fraudulent leads to Merchant2 and earning commission from them. Upon internal investigation by Merchant2's fraud team, Merchant2 had decided not to honour the commission earned by Stan Smith. Merchant2 then raised an incident report with AffiliateNow to warn other merchants of his fraudulent behaviour and to have him removed from the network.

The AffiliateNow employee investigating Stan Smith's case found that the traffic being sent to merchants from Stan Smith's affiliate account was coming from the same referring site and many of the IP addresses were repeated. The repeated IP addresses were all from foreign countries and visiting sites that only offered services in the UK. AffiliateNow suspended Stan Smith's account and issued a warning to all affected merchants about Stan Smith's account.

That is as far as the incident report we have received seems to go. However, upon further investigation, we have found several links from Stan Smith to other accounts in the affiliate database. The originally reported account is linked to 5 other affiliate accounts in the database, four of which are listed in Stan Smith's name with the 5th having his name in the cheque payable field and the name Edward Smith as the account holder. These different accounts have four unique physical addresses, two of which have been listed by other affiliates as their address.

Three of the accounts we examined have different names in the cheque payable and account holder fields. Of the five affiliates with bank account information on file, three also listed a different name on the bank account than that of the account holder.

The greatest anomaly that we discovered involved the telephone numbers listed for the affiliates.

Surprisingly, only six out of the 28 affiliates examined listed mobile phones when asked for a phone number. This is helpful to our analysis because landline telephone numbers can be traced back to a general area. Only one of the remaining 22 numbers, however, had a dialling code consistent with the address information provided by the affiliate. This should be a significant clue that something is amiss with the accounts of these affiliates.

6 Proposed Detection Methods for Affiliate Scams

In looking at the customer database, several inconsistencies are present. A system designed to seek out these anomalies may enable affiliate networks to flag accounts for a closer inspection by an employee. For example, if a detection system had been run in our case study it may have picked up that Stan Smith was registered to multiple physical addresses. It may have also picked up that multiple users were registered to these addresses as well. Linking these accounts together may enable the affiliate network to remove large chunks of fraudulent accounts with a single investigation rather than a new investigation for each account.

We believe that the most suspicious detail in the customer records is the fact that none of the telephone numbers originate from the area listed in the address details of the affiliate. A person listed as living in Gotham City may have a phone number with a dialling code for Shelbyville, for example. A comparison between the dialling code and postcode of the listed address could easily mark such an account at a high risk for being fraudulent if a landline phone number is provided during the affiliate registration process.

Another telltale sign of fraudulent behaviour is an in-depth look at the affiliate's site. If the site consists mainly of banners and ads, or is in some other way inappropriate for the products being advertised, the page may belong to a malicious affiliate. If the affiliate does not use proper grammar and complete sentences it may be a sign that the site was hastily made. If the images appear broken or are taken from another site, something suspicious may also be going on with the affiliate. These are a couple of the more obvious signs that the site may have been created simply to host the ads and scripts necessary to generate a fraudulent income.

Less obvious signs might be found in the code of the affiliate's site. Websites that contain scripts used by known fraudsters such as the code example shown in Figure 2 should probably be looked at more closely. If an affiliate is producing dozens of sites for their operation, they are likely to all have a similar layout and similar mistakes in their code. Running the website through a HTML and CSS validation checker on suspected pages may produce similar results, which could be an indication of multiple accounts involved in dodgy behaviour.

Weighting each of these categories and keeping track of an affiliate's score while running these tests could give an indication of whether or not the affiliate is genuine, fraudulent or undetermined. In the case of fraudulent and undetermined, the case could be moved to the fraud team of the affiliate network for further investigation.

Apart from the affiliate database and sites, a good indication that an account is involved with fraudulent behaviour is duplicate IP addresses appearing from the same affiliate on multiple merchants. An occasional duplicate IP address is not necessarily fraudulent, but the same duplicate IP address(es) multiple times in a small time period is pretty suspicious.

Another method of combating the rising number of malicious affiliates is to prevent them from joining a programme in the first place. Edelman posits that it may be possible to prevent fraudsters from joining an affiliate programme all together (Edelman 2007b). He found that if a merchant pays their affiliates in arrears with compensation to offset the extra time before payment is received, there exists a certain point at which it is no longer profitable for fraudsters, or bad-type agents as he calls them, to participate in the programme. Unfortunately, according to a recent survey (AffStat 2009) of over 450 affiliates, 57% of good-type affiliates decide whether to join an affiliate programme based upon how often a programme pays out. With the majority of affiliates basing their preference of programme on how soon they start earning, extending that wait may decrease the number of good affiliates a merchant or affiliate network can attract.

7 Conclusions and Future Work

To help overcome the affiliate network scam, a key component is the methods of analysing affiliate's details and then ranks that affiliate based upon the likeliness of the account being created for fraudulent purposes. This could save companies time and money by removing malicious affiliates before they have had a chance to commit any fraudulent acts and labelling affiliates that should be watched more closely for indications of fraudulent activities.

The current research aims to provide an anonymisation framework for investigations, and has presented an example using a real-life affiliate scam. While using a blanket anonymisation technique has caused us to lose a slight amount of context when dealing with location data, we feel that in this case, the loss did not hinder our ability to find connections between affiliate accounts or to draw conclusions about the involvement of the affiliates examined.

The next phase of the research involves building the anonymisation framework, which will feed into an analysis engine. The analysis engine will check the both the affiliate's profile and build upon the system developed in (Edelman 2007a) to check the affiliate's websites for known attacks. A metric based upon the *quality* of the affiliate's web site may also be useful to alert merchants and networks to the presence of malicious affiliates that may be involved in fraudulent behaviour.

8 References

7safe & ACPO (2007), Good practice guide for computer-based electronic evidence, Technical report, ACPO and 7Safe.

AffStat (2009), 2009 affiliate summit affstat report, Affiliate marketing benchmarks, AffStat.

Bloch, M. & Eroshenko, D. (2004), How to defend your website against click fraud, White paper, Clicklab.

Ciriani, V., De Capitani Di Vimercati, S., Foresti, S. & Samarati, P. (2007), *k-anonymity*, volume 33 edn, Springer US, pp. 323–353.

Edelman, B. (2007*a*), 'Introducing the automatic spyware advertising tester', Online. http://www.benedelman.org/news/052107-2.html

Edelman, B. (2007*b*), 'Spyware still cheating merchants and legitimate affiliates'. http://www.benedelman.org/news/052107-1.html

Edelman, B. (2008*a*), 'Auditing spyware advertising fraud: Wasted spending at vistaprint', Online. http://www.benedelman.org/news/093008-1.html

Edelman, B. (2008*b*), 'Cpa advertising fraud: Forced clicks and invisible windows', Online. http://www.benedelman.org/news/100708-1.html

Edgar, D. (2004), Data sanitization techniques, White paper, Net 2000 Ltd.

Google (2009), 'Google adsense program policies', Online. https://www.google.com/adsense/support/bin/answer.py?answer=48182

Grow, B., Elgin, B. & Herbst, M. (2006), 'Click fraud: The dark side of online advertising', Online. http://www.businessweek.com/magazine/content/06_40/b4003001.htm

Holton, K. (2009), 'UK internet ad spend overtakes TV for first time', Online. http://www.reuters.com/article/idUSTRE58S4IL20090929

Metwally, A., Agrawal, D. & El Abbadi, A. (2005), Duplicate detection in click streams, *in* 'WWW '05: Proceedings of the 14th international conference on World Wide Web', ACM, New York, NY, USA, pp. 12–21.

Ntoulas, A., Najork, M., Manasse, M. & Fetterly, D. (2006), Detecting spam web pages through content analysis, *in* 'WWW '06: Proceedings of the 15th international conference on World Wide Web', ACM, New York, NY, USA, pp. 83–92.

Tuzhilin, A. (2006), The lane's gifts v. google report, Technical report, Stern School of Business at New York University. http://www.reference.com/go/http://googleblog.blogspot.com/pdf/Tuzhilin_Report.pdf

Tveit, A., Edsberg, O., Røst, T. B., Faxvaag, A., Øystein Nytrø, Nordgård, T., Ranang, M. T. & Grimsmo, A. (2004), Anonymization of general practitioner's patient records, *in* 'Proceed-

ings of the HelsIT'04 Conference'.

Venkatesan, T., Gupta, H., Roy, P. & Mohania, M. (2008), 'Efficient Techniques for Document Sanitization', *pages.cs.wisc.edu* . http://pages.cs.wisc.edu/venkat/docsan.pdf

Zango (2005), '180solutions sues former affiliates for illegal software installations'. http://www.zango.com/Desintation/Corporate/ReadArticle.aspx?id=29

Zhang, L. & Guan, Y. (2008), Detecting click fraud in pay-per-click streams of online advertising networks, *in* 'ICDCS '08: Proceedings of the 2008 The 28th International Conference on Distributed Computing Systems', IEEE Computer Society, Washington, DC, USA, pp. 77–84.

Zhong, S., Yang, Z. & Wright, R. (2005), Privacy-enhancing k-anonymization of customer data, *in* 'Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems', ACM, p. 147. http://portal.acm.org/citation.cfm?id=1065167.1065185

Anon_id	Tel Area	Anon_First	Anon_Last	Snailmail	Snailmail2	Logins	Earned	Registered
100026	Gotham City (A)	Stan	Smith	1 Spooner Street	Gotham City (G)	13	£0.00	26/12/2005
100027	Gotham City (A)	Edward	Smith	1 Brookside Close	Gotham City (J)	41	£0.00	12/08/2005
100006	Mobile	Stan	Smith	10 Evergreen Terrace	Gotham City (D)	206	£801.00	18/06/2006
100003	Gotham City (B)	Robert	Johnson	10 Evergreen Terrace	Gotham City (D)	183	£1,180.00	15/10/2006
100004	Liberty City	Robert	Johnson	10 Evergreen Terrace	Gotham City (D)	416	£8,103.10	20/03/2007
100007	Sunnydale	Frank	Smith	10 Evergreen Terrace	Gotham City (D)	232	£2,404.25	11/04/2007
100008	Shelbyville	Richard	Miller	10 Evergreen Terrace	Gotham City (D)	15	£535.00	30/09/2007
100002	Springfield	James	Williams	10 Evergreen Terrace	Gotham City (D)	10	£150.00	12/10/2007
100010	Gotham City (B)	Joseph	Garcia	10 Evergreen Terrace	Gotham City (D)	11	£285.00	17/10/2007
100013	Gotham City (C)	Daniel	Taylor	10 Evergreen Terrace	Gotham City (D)	147	£3,781.96	24/10/2007
100014	Mobile	Paul	Martin	10 Evergreen Terrace	Gotham City (D)	3	£0.00	06/03/2008
100011	Petoria	Thomas	Anderson	10 Evergreen Terrace	Gotham City (D)	166	£1,052.00	15/03/2008
100009	Mobile	Charles	Davis	10 Evergreen Terrace	Gotham City (D)	94	£28.00	09/04/2008
100012	Gotham City (I)	Christopher	Anderson	10 Evergreen Terrace	Gotham City (D)	220	£50.00	17/04/2008
100015	Gotham City (D)	Mark	White	10 Evergreen Terrace	Gotham City (D)	13	£0.00	15/05/2008
100001	Mobile	John	Doe	10 Evergreen Terrace	Gotham City (D)	7	£0.00	23/07/2008
100005	Ogdenville	Craig	Smith	10 Evergreen Terrace	Gotham City (D)	56	£69.00	27/07/2008
100000	Mobile	Mike	Rogers	10 Evergreen Terrace	Gotham City (D)	17	£0.00	02/08/2008
100022	Gotham City (E)	Christopher	Anderson	416 Cherry Street	Gotham City (C)	316	£1,602.70	10/07/2006
100025	Gotham City (F)	Aaron	Robinson	416 Cherry Street	Gotham City (C)	45	£45.00	29/04/2007
100023	Gotham City (G)	Caleb	Harris	416 Cherry Street	Gotham City (C)	18	£480.00	28/09/2007
100016	North Haverbrook	Matthew	Brown	416 Cherry Street	Gotham City (C)	14	£220.00	02/10/2007
100021	Gotham City (B)	Nicholas	Wilson	416 Cherry Street	Gotham City (C)	13	£300.00	07/10/2007
100018	Mobile	Stan	Smith	416 Cherry Street	Gotham City (C)	267	£587.00	14/11/2007
100017		Tony	Jones	416 Cherry Street	Gotham City (C)	74	£132.24	26/12/2007
100024	Gotham City (F)	Stan	Smith	416 Cherry Street	Gotham City (C)	66	£318.10	07/03/2008
100019	Gotham City (H)	Craig	Smith	416 Cherry Street	Gotham City (C)	16	£128.00	31/03/2008
100020	Protected	Craig	Smith	416 Cherry Street	Gotham City (C)	9	£0.00	04/04/2008

Table 1 - Sample of Anonymised Data