

A black market for upvotes and likes

This article investigates controversial online marketing techniques that involve buying hundreds or even thousands of fake social media items, such as likes on Facebook, Twitter and Instagram followers, Reddit upvotes, mailing list subscriptions, and YouTube subscribers and likes. The findings presented here are based on an analysis of 7,426 “campaigns” posted on the crowdsourcing platform microworkers.com over a 365 day (i.e., year-long) period. These campaigns contained a combined 1,856,316 microtasks with a net budget of USD 208,466.

Keywords: *Crowdsourcing, facebook, twitter, youtube, reddit*

Author Information

Mihaly Heder, Budapest University of Technology and Economics
<https://orcid.org/0000-0002-9979-9101>

How to cite this article

Heder, Mihaly. „A black market for upvotes and likes”,
Információs Társadalom XIX, 4. no (2019): 18–39.

<https://dx.doi.org/10.22503/inftars.XIX.2019.4.2>

*All materials
published in this journal are licenced
as CC-by-nc-nd 4.0.*

A black market for upvotes and likes

Mihály Héder

Abstract

This article investigates controversial online marketing techniques that involve buying hundreds or even thousands of fake social media items, such as likes on Facebook, Twitter and Instagram followers, Reddit upvotes, mailing list subscriptions, and YouTube subscribers and likes. The findings presented here are based on an analysis of 7,426 “campaigns” posted on the crowdsourcing platform microworkers.com over a 365 day (i.e., year-long) period. These campaigns contained a combined 1,856,316 microtasks with a net budget of USD 208,466.

Keywords: Crowdsourcing, facebook, twitter, youtube, reddit

Introduction

This article analyses marketing campaigns that have been executed through the hiring of *freelancers* or “*microworkers*” to complete short, menial tasks called *microtasks* that usually pay less than one US dollar each, and most often only around ten cents. These tasks include watching, liking, upvoting, and “+1”-ing items on web platforms featuring social media functions, like Facebook, Twitter, Reddit, and Instagram. A job description showing what such a campaign looks is given below (an actual, observed example):

Title: Facebook Like: <REDACTED String>

Payment: USD 0.15

Number of workers accepted: 100

Job description:

WARNING: we manually review almost all of the submitted tasks. Thus, if we find that you have ignored the instructions (i.e., posting on a wrong site, using non-unique or nonsensical content), you will be permanently banned from our system. You must have 50 Facebook friends.

1. Go to <REDACTED URL1>
2. Visit the URL shown at <REDACTED URL1>
3. On that page, you will see a Facebook Like button. Click on it
4. Submit your Facebook profile URL on <REDACTED URL1> to verify that you have completed the task (make sure you have set your Facebook profile to Public View in order for your task to be verified) to get a 7 character confirmation code (...)

Campaign example 1: Buying facebook likes

Completing this task pays USD 0.15 for the microworkers. In this case, one hundred freelancers were sought to do the task.

There are specialized web platforms for the brokering of small tasks. One of them is microworkers.com, which was investigated for this article. This platform lets employers run campaigns for a fee, usually 10% of the payment made to the freelancing internet users who do the job, called microworkers. Therefore, the 100 Likes above would cost the client posting the job only USD 16.5. The platform has hundreds of thousands of microworkers (Nguyen 2014).

The platform was not specifically created for promotional or marketing purposes, and indeed data processing, survey-based research, and software testing jobs are also posted on the platform. However, as shown below, the majority of tasks can be categorized as being of a promotional nature (See Table 3).

1.1. Article Structure

The structure of the present article is as follows. After this introduction, the terminology is presented and a number of ethical questions are posed. This is followed by an outline of the related work in this field. The next chapter describes the research aims and details of the observation, as well as the limitations of the research. This is followed by the Results section, which also contains a number of separate subchapters covering various aspects of the investigation into the most important gray promotional activities, including estimates about their efficiency, security concerns, and possible counter-measures, where applicable. The Results section also describes some observed campaign-supplementing techniques. Finally, the Conclusions section offers a summary and outlines some possibilities for future work.

2. Why call it a black market?

In the opinion of the author of this paper buying likes, followers, votes, upvotes, retweets, etc. (the generic term for these used in this paper is social media activity) for promotional purposes is an unethical practice, therefore we use the term “black market” to describe this part of the industry. This is not to say they are necessarily breaking any rules or laws—that question is outside the scope of this paper.

This kind of activity is considered unethical for the following reasons. a) The users of a social platform are misled by a page or post having an artificially inflated number of likes, followers, etc. Normally, it is impossible for users to differentiate between paid activity and genuine activity. The conventional semantic behind a like, upvote, or follow is a statement of approval of the content in question. In other words, the microworkers are paid to lie in the sense that they are paid to pretend to like/endorse/approve content that they most likely have not truly read or watched. b) The microworker’s “employment” is arguably rather exploitative, because of the low payment and the apparent lack of any powers against the clients (see table 5). c) Content creators and social media users who do not employ such practices are clearly disadvantaged. Finally, d) it is easy to see that if paid social media activity were more universal, it would create an unsustainable social media environment.

Besides paid social activity, there are other highly controversial campaigns based around social media. One such typical activity involves creating accounts on behalf of a

client and then the microworker handing over the user name and password to the client. An example of this kind of activity is given in the following (actual, observed example):

Title: Gmail: Create an Account
Payment: 0.16

Number of workers accepted: 230

Job description:

Note: I will Check American Name and Profile Picture otherwise I have to decline you.

1. Go to www.fakenamegenerator.com
2. Choose proper American name
3. Go to Gmail.com
4. Create a new Gmail account using details from www.fakenamegenerator.com
5. Upload a good profile picture in Gmail

For proof:

- Give Password
- Give a Recovery Mail so that I can change it later

Note: Make sure all are perfect otherwise, I will decline your payment.

Required proof that task was finished?

1. Gmail and Username
2. Password
3. Recovery email

Campaign example 2: Acquiring gmail accounts

In this case, the client was able to acquire 230 Gmail accounts for a mere \$ 40.48. We might speculate that these accounts (and similarly those created for YouTube, Twitter etc.) will be used for promotional purposes, controlled by the client, while also raising all the ethical concerns highlighted in the previous example. However, this speculation might actually be optimistic. Fake accounts like these could also be used for more sinister purposes, such as as tools for spreading fake news (Allcott und Gentzkow 2017) for political purposes or for committing fraud.

To sum up, it seems justified to state that there is a black market for “fake” accounts and social media activity and that promotion done using this market represents an ethically gray area.

The author wants to point out that this is not to say that microworkers.com or any other platform is in itself immoral or designed to be a black market—such platforms also offer a useful venue for valid projects, like data processing at scale, acquiring subjects for survey-based or interactive online scientific research, monitoring a competitor’s public online activity, or counting objects on images - several such campaigns were observed.

3. Related Work

The marketing techniques discussed in this article are related to techniques called sock puppetry and click farming. Sock puppetry means the control of many social media accounts by one person or a small group of individuals. A report on such activity was published in The

New York Times (Caldwell 2007). The method of control is a crucial difference in these efforts. As we will see some in Section 6.5 clients buy hundreds or thousands of accounts that they can use themselves. In this case there are technical possibilities to detect the puppetry, by noticing when a very high number of users are logging in from the same network location or use the browser client fingerprint (Laperdrix u. a. 2016).

But the method of control can also be an order from a client to a cohort of users to perform some activity (but the client itself never logs in to any of their accounts). In this case detection of the activity is much harder if the client takes some precautionary steps (see Section 6.5). Sometimes this is called meet puppetry (Cook u. a. 2014) referring to the fact that it involves real freelancers.

Click farming is another related term. Click farms are actual workplaces in developing countries where a large number of employees are performing short task sometimes for as low as 1000 likes for 15 USD (Arthur 2013). In current reporting these are often called troll farms (Smith 2018).

The platform investigated by this article, microworkers.com differs greatly from a click farm as it is a completely distributed crowdsourcing tool, but some of the campaigns done here might be similar to those done by a click farm.

On the deceptiveness of such campaigns in comparison with traditional advertising (where the prospective customer is aware that it is being presented with an ad) is well described by Del Riego (2009) and Forrest und Cao (2010) in connection with then-new US Federal Trade Commission guidelines on endorsements and reviews.

This article focuses on probably way smaller market available on microworkers.com, which is, however, easily accessible for freelancers and is not specialized to gray marketing in particular. In contrast with the services that directly offer followers and social media activity (Fiverr, SeoClerks, InterTwitter, FanMeNow, LikedSocial, SocialPresence, SocializeUk, ViralMediaBoost (De Micheli und Stroppa 2013)), here the client has to organize its own campaign and orchestrate the freelancers on the crowdsourcing platform. This allows for creativity and innovations in the campaign methods.

Nguyen (2014) explained the idea behind microworkers.com, founded in 2009, as well as reported its user count at the time of writing (presumably 2014). Howe (2006) also reported on microworkers.com as a crowdsourcing platform.

According to Nguyen, the platform had over 600,000 users from 190 different countries. The aim of the microworkers.com as a project was to aid brokering crowdsourcing campaigns. As the article explained,

In crowdsourcing platforms, there is perfect meritocracy. Especially in systems like Microworkers; age, gender, race, education, and job history do not matter, as the quality of work is all that counts; and every task is available to Users of every imaginable background. If you are capable of completing the required Microtask, you've got the job.

The campaign templates on the landing page of the platform are great sources of inspiration for what could possibly be achieved through crowdsourcing: participating in market research, captioning documents and video, categorizing images, testing websites and applications, and so on. The fact that the majority of public campaigns visible on the platform are mostly employing controversial promotion techniques does not seem to be the result of the platform design or intentions.

Hirth u. a. (2011) investigated microworkers.com in order to compare it to the much better understood Amazon Mechanical Turk (Paolacci u. a. 2010, Buhrmester u. a. 2011). They correctly identified a main difference between the portals: the payment mechanism. At the time, it was basically impossible to use MTurk without a US-based credit card, while the microworkers website allowed payments to be made with Moneybookers (called Skrill today). This helps explain why the author of this paper and possibly other non-US-based researchers first discovered microworkers. Works by Gardlo u. a. (2012) and Crone und Williams (2017) aimed to assess the usefulness of the platform for scientific purposes; and indeed, scientific projects regularly, though relatively infrequently, appeared on the platform. However, it is possible that this difference in payment methods is only one of the reasons behind the nature of the campaigns conducted on each.

Hirth et al.'s work (Hirth u. a. 2011) indicated that gray promotional campaigns were already existed as long ago as 2011: "Signup", "Click or Search", "Voting and Rating" were already featured as campaign categories; however, the payments offered were slightly higher than today.

The connection between social media and marketing was analyzed by Thackeray u. a. (2008) as early as 2008. Of course this work concentrated on the legitimate social media strategies firms might embrace, such as paid search results, where the brand buys a presence in the search results. As Yang und Ghose (2010) explained, these are usually placed in a separate area on the results page, together with being clearly to indicate that they are paid for or they may be labeled as an ad, e.g., on Facebook (the difference between "paid" and "organic" (showing up in non-paid results) links is less emphasized in today's search engines but remains clear). Rutz und Bucklin (2011) demonstrated how tying paid search results to generic search terms might increase the success of a branded paid search.

The search and engage campaigns (see section 6.2) discussed in the present paper are different. They don't try to increase visibility by buying paid results. These represent the dark side of search engine-based advertising, whereby they try to directly manipulate the organic links.

It was envisaged (Zhang u. a. 2013) that the identification of the key influencers on social media could be crucial for effective viral marketing—but with the gray marketing techniques presented here, influence is attempted to be created directly, albeit artificially.

It was also envisioned that customer-generated content on blogs, etc. would be crucial for promotional activities—in the present paper, campaigns seeking to manufacture legitimate-looking customer content are analyzed. In other words, these campaigns, albeit unethical, are sometimes the effective counterparts of hard-to-operate social media marketing tactics, or in other words, they are controversial shortcuts to followers and likes and brand-friendly social content.

Confessore et al.'s recent work reported on in the New York Times (Nicholas Confessore und Hansen 2018) covered a very similar theme to this article, but was focused mostly on Twitter and on the activities of a company called Devumi.

4 Aims and Methods

4.2 *Research aims*

The primary aims of the present research were to identify the different schemes employed in microworker-based gray marketing on microworkers.com and to then categorize them,

attempt to uncover how they fit in a wider strategy, estimate their limits and effectiveness, and to utilize this knowledge to provide general insights into these kinds of campaigns.

Second, the scale and typical budget of these campaigns, and their relative share of the overall activity on the microworkers.com platform were also measured and reported herein.

4.2 Observation of campaigns

This research project attempted to observe all campaigns posted on microworkers.com from 22 February 2016 to 21 February 2017. The site was checked several times a day during this period. In total, 7,426 campaigns were observed during the period. Each campaign was manually categorized and the aggregate numbers of categories (payments, number of tasks) were updated.

4.3 Categorization

The campaigns were categorized in terms of two dimensions: the related target platform (Facebook, Google Plus, SoundCloud, Twitter, etc.), and the specific activity (search and engage, like, comment, sign up, etc.). The categorization was based on the campaign title and text and proved to be quite straightforward as the names of the platforms were clearly stated in the title and represented unambiguous brand names, and as the activity to be performed was almost always explained in an itemized list in the posting. Obviously certain activities were further linked to certain platforms (like retweets can only be done through Twitter), but others, like search and engage can be done on multiple platforms.

The platform labels utilized are summarized in Table 1. while the activity labels are summarized in Table 2.

Also, each campaign could belong to multiple activity categories, by using multiple tags.

Platform categories

Ali (Alibaba, AliExpresS)	Instagram
Amazon	Bing
MBS: (microblog instant share): blog- ger.com, Pinterest, Digg, Tumblr, 9gag or other blogs or quick sharing platforms	Mix: (SoundCloud, Mixtape, datpiff)
Browser add-on (e.g., Chrome extensions)	RDT (a traffic generator site, the name of which is redacted from this article)
Other: (500px, Wordpress.com, Snapchat, Skillshare, Hotmail, LinkedIn, Coursera, Bitbucket, Snapchat, Steam, Yandex, other uncategorized)	Question sites: (Yahoo Answers, Quora)
eBay	Reddit
Forum (Disqus, Warrior Forum, other forums)	Redacted: (not visible in description because of the rotator technique, explained below)
Facebook	Smartphone (iOS and Android)
Gmail	Twitter
Google (search)	Yahoo (search)
Google+	

4.4 Limitations

On microworkers.com, there are invite-only campaigns as well. Unfortunately, there is no information freely available on these campaigns or their share of the total number of campaigns. The nature of invite-only campaigns, for which clients apparently hire tested and trusted microworkers, could be a subject for further research.

Some campaigns might have slipped through between two observations, meaning that their full life cycle very short (only a couple of hours). This does not appear to be typical but cannot be ruled out. Therefore, it can be said that in reality there were possibly more than 7,426 public campaigns and an additional, unknown number of invite-only ones.

As explained above, the campaigns were manually categorized by the author. This categorization, because it relied on objectively observable features of the campaigns, did not require significant subjective judgment. Therefore, in the context of intersubjectivity, it should not be a serious limitation that there was no multiple-person cross-checking performed for interpretation of the category labels.

Finally, there are obviously other brokering platforms for such microtasks, but these are outside the scope of this investigation (in fact, some of those platforms seem to be using microworkers for recruitment purposes).

The author is confident that these limitations do not prevent the work from meeting its stated aims, as it is an explorative rather than exhaustive description of techniques and strategies.

4.5 Anonymization

Generally, all data presented in this paper (mostly microtask descriptions) is anonymized. Most of the actual URLs, person and company names, and other named entities are replaced with the string `<REDACTED(...)>`. When there are several URLs or names within

Activity categories

Ali (Alibaba, AliExpresS)	Instagram
Amazon	Bing
MBS: (microblog instant share): blog- ger.com, Pinterest, Digg, Tumblr, 9gag or other blogs or quick sharing platforms	Mix: (SoundCloud, Mixtape, datpiff)
Browser add-on (e.g., Chrome extensions)	RDT (a traffic generator site, the name of which is redacted from this article)
Other: (500px, Wordpress.com, Snapchat, Skillshare, Hotmail, LinkedIn, Coursera, Bitbucket, Snapchat, Steam, Yandex, other uncategorized)	Question sites: (Yahoo Answers, Quora)
eBay	Reddit
Forum (Disqus, Warrior Forum, other forums)	Redacted: (not visible in description because of the rotator technique, explained below)
Facebook	Smartphone (iOS and Android)
Gmail	Twitter
Google (search)	Yahoo (search)
Google+	

Table 2: Categories of activities

one example though, they are replaced with their own unique label so that they are not mixed up. Other than this modification, the job descriptions are copied herein verbatim.

Obviously, some basic URLs like Facebook.com or fakenamgenerator.com, for example, are kept because they are reported only for uncovering the campaign method but not its content, and also because the job descriptions would not be as understandable without them.

The reason for the redaction is that the persons, websites, and Facebook accounts mentioned in these task descriptions may be unwilling targets of a campaign. It is also probable that the customers of promotion campaigns are often not aware or may even have been misled about the methods employed on their behalf.

5 Results

Based on the observations, the following summaries were created.

Table 3 contains budget summaries by activity. The columns represent the activity code, number of campaigns, number of tasks, and net budget (without the 10% fee). The table is ordered by the descending number of campaigns.

Activity	# campaigns	# tasks	total budget
L	1,303	207,811	\$ 22,757.32
P	1,293	116,682	\$ 26,796.95
S	1,229	577,444	\$ 46,802.91
U	1,210	354,180	\$ 40,344.71
C	733	227,756	\$ 27,926.15
E	495	318,387	\$ 23,310.10
I	361	16,934	\$ 8,375.01
H	357	26,305	\$ 7,547.54
Z	352	122,974	\$ 9,681.65
N	203	26,808	\$ 3,419.92
B	196	40,045	\$ 4,422.94
W	147	38,756	\$ 4,307.35
D	138	50,886	\$ 6,307.99
F	79	12,649	\$ 1,520.29
V	68	32,832	\$ 4,026.03
T	31	2,071	\$ 409.87
R	29	5,191	\$ 2,037.10
O	24	7,219	\$ 1,240.66
A	14	519	\$ 78.90
M	12	847	\$ 203.45
K	10	474	\$ 292.02
Q	3	150	\$ 24.60
Y	1	1,026	\$ 61.56

Table 3: Budget Summary by Activities

If we take out the category data processing, research participation, captcha solving, testing, installing and “other”, the remaining activities are purely for promotional purposes. This leaves 1,665,138 tasks, or 89.7% of the whole. It should be added that many of the

install tasks appear to be promotional (see section 6.4). Counting these in would make the figure even higher. However, since for many such campaigns this aspect is impossible to tell, they are left out.

A similar summary for the platforms is given in table 4. The first column is the platform name, the rest is the same as before. The table is ordered by the descending number of campaigns.

From these values, the average payment for tasks related to certain activities and platforms could be calculated. Here is the distribution of the payments (not equal ranges):

The lowest paid wage was \$0, and this was incidentally the biggest campaign with 99,999 tasks. The task was a simple visit to a link. The client promised future tasks for those who completed the task. More detail is provided on this task in section 6.5.

The highest paid wage was \$3 for a task, but as can be seen from the above table, there were only 531 jobs offering between \$1.1 and \$3, while there were over a million jobs offering between 5 and 10 cents, making the higher paid tasks very rare indeed.

6 Campaign types and their analysis

6.1 Voting

Participation in voting (V) is a recurring activity on microworkers.com, with 68 campaigns posted featuring an aggregated 32,832 votes purchased. The top two voting campaigns seemed to be promoting a product and a sports team (2,130 and 2,000 tasks), the third was

Platform	# campaigns	# tasks	# total budget
Smartphone	1102	231,892	\$ 27,702.07
Google+	823	57,203	\$ 18,960.50
Other	797	349,937	\$ 31,333.06
Twitter	747	83,731	\$ 14,047.21
Facebook	666	115,967	\$ 15,657.86
YouTube	616	248,914	\$ 29,322.28
Reddit	539	77,104	\$ 5,434.15
Redacted	501	132,401	\$ 13,397.94
Google	330	261,923	\$ 18,672.37
RDT	296	80,844	\$ 8,034.94
MBS	218	38,859	\$ 5,497.36
Instagram	172	19,421	\$ 2,072.35
Amazon	156	55,760	\$ 6,934.85
Mix	144	11,803	\$ 1,458.55
Question site	126	16,190	\$ 2,447.59
Forum	84	7,558	\$ 964.15
Gmail	50	10,275	\$ 2,908.35
eBay	22	6,014	\$ 570.52
Yahoo	21	45,686	\$ 2,355.13
Ali	9	3,945	\$ 551.91
Bing	4	675	\$ 69.00
Browser add-on	3	214	\$ 73.96

Table 4: Budget Summary by Platforms

a giveaway voting for an expensive trip for couples, where the entrants were supposed to vote on the videos they uploaded about themselves. One entrant purchased 1,703 votes (the wording reveals that they bought the vote for themselves personally) for \$0.12 each:

Title: Video: vote

Payment: 0.12

Number of workers accepted: 1703

1. **Go to <REDACTED URL>**

Give the video a VOTE by clicking on the heart below the video Required proof that task was finished?

1. **Tell me how many votes I had after you've voted**

Campaign example 3: Buying votes

Unfortunately while this can be seen as being clearly unethical, for a little more than \$ 200 it could have been economically viable if the entrant won the vote. And it is even plausible that a local vote could have been won with just 1,703 extra votes bought.

Other votes were for titles like best auto repair shop, best bakery or “tradie of the year in Australia”. There was also a census on how many New York City residents wanted to go on a date with a certain model. There were votes on temple photography, the best fintech firms, music mixes, the best female vocalists, several contests about the ranking of attractive persons, a vote on XXL Magazine, a vote on the best local charter flight provider, and so on. Some of these were clearly promoting a product or a performer or artist; others seemed to be clearly what we will call vanity-promotion.

The purchased votes ranged from dozens to about 1,000 at a time. It is hard to assess the overall efficiency of such campaigns, but it can certainly be said that for local contests, where the maximum number of people voting is expected to be measured in hundreds or thousands, it is very easy to rig contests this way, as it would cost only a few dollars.

task payment	total number of tasks
\$0	99,999
\$0.05-\$0.1	1,059,172
\$0.11-\$0.2	581,146
\$0.21-\$0.3	63,904
\$0.31-\$0.5	35,893
\$0.51-\$1.0	15,671
\$1.1-\$3.0	531

Table 5: Number of tasks by payment range

6.2 Search and engage tasks

The common feature of these kinds of microtasks seems to be an attempt to manipulate the search results in search engines like Google, Yahoo, Bing or the search feature of Facebook. In most cases, a certain item is promoted, but in some rare cases, the goal actually seems to be to push unwanted result items back in the result list.

These campaigns appear to assume that search engines learn: if for given search term, a high number of users click on a particular result item, then that result item must be a good result for the search and therefore it will be listed early in the results listing. While the actual algorithms search engines use are proprietary and unpublished, it is known how they work in theory (Buttcher u. a. 2016). It is thus plausible that they can be tricked in this way to a certain extent. We know that user behavior is taken into account in Google, for instance, as (Clark 2015) reported that Google’s novel AI solution, BrainRank, was the third most important factor (the technical term is “signal”) when ranking pages. We also know that it learns from user behavior, hence it is plausible these systems can be tricked through paid user behavior simulating genuine interest.

A search and engage campaign therefore hires a large number of microworkers for searching the given terms and clicking on the promoted item in the search results.

An example of this type of campaign is given below:

Title: Google search

Payment: 0.08

Number of workers accepted: 1600 **Job description:**

1. Open up Google.com (please use US version)
2. In Google, please search for this phrase: <REDACTED PERSON NAME>
3. Please click on any of the red boxed links you see in the attached file <the file is a Google result list screenshot, red rectangles designate what result items need to be promoted>
4. Stay on page for 1 minute (...)

Campaign example 4: Google search

The top ten search and engage campaigns have the number of tasks offered as between 2,941 and 6,770. However, many of these campaigns seem to be part of the same project, making the biggest projects around 10,000-20,000 tasks.

An interesting tendency is that in many of these promotion campaigns, it seems that there is no marketed product involved, rather it is individuals concerned with their online persona who are the payers, in what is really another example of vanity-promotion.

The biggest search and engage project, with above 10,000 tasks, involved the promotion of a USA business executive’s Wikipedia entry, whose name happens to be the same as a famous USA American football player and also a former USA congressman. The project must have been a success as currently the promoted page comes out top in a Google search when searching for that name. Naturally, it is impossible to establish the causal relation between the campaign and the current ranking with any certainty, especially this long after the campaign.

6.3 Social media activity

Paid social media activity involves task like creating Pinterest Pins, upvoting in Reddit, YouTube, or Google+, liking in Facebook, using Digg, Twitter, or Instagram, commenting on forums, upvoting on SoundCloud, Mixcloud or Datpiff, and so on. In terms of activity codes, this section covers B, C, F, H, L, N, T, W.

Campaign example 1 is in this category. The campaigns are usually straightforward and easy to do, therefore the payment is usually very low. Clicking on like, upvote, etc. are the lowest paid tasks. For instance, the 77,104 Reddit upvotes purchased during the 365 days study period cost less than \$ 5,500 in total (see Table 4). The highest paying jobs in this category were those that require writing content that meets a set specification, e.g.:

Title: YouTube: Comment 3x (1-3)

Payment: 0.30

Number of workers accepted: 90 **Job description:**

1. Go to www.youtube.com/channel/<REDACTED>.
2. Post a positive relevant comments on the videos found in first three links Important: Comment must be at least 10-15 words long and cannot be generic and must include the following words <REDACTED Person name> and the word "Florida".
3. Stay on each video page for 1 minute

Required proof that task was finished?

1. YouTube display name
2. Copy of the comments you've posted
3. URL to YouTube videos where comments were posted

Campaign example 5: Youtube comments

However, in other cases the freelancer is asked to copy-paste the comment content, and the job then pays less:

Title: YouTube: Comment 3x (<REDACTED>)

Payment: 0.12

Number of workers accepted: 440 **Job description:**

1. Go to the instruction page: <REDACTED URL>
2. Search Youtube.com for the key phrase
3. Copy-paste the supplied comments onto relevant video
4. Repeat steps 2 and 3 for 2 more videos (3 total)

Required proof that task was finished?

1. Your YouTube name
2. The search phrases you used
3. URLs of the 3 videos you commented on

Campaign example 6: Youtube comments - the copy-paste method

For commenting, the most prominent platforms are YouTube (364 campaigns; 197,735 tasks, some paying for three comments), Instagram (101; 4,670), questions sites (81; 10,282), Facebook (30, 2,485), and a long tail of other forums (see some under Platforms; 118 campaigns).

Following a given account is done on Google+ (317 campaigns; 21,228 tasks) Instagram (54; 8,960), Twitter (20; 2325), Quora and Yahoo Answers (4; 466), and Google+ (1; 898). It must be noted that for Twitter, the clients often require the presence of some features from the freelancer, e.g.:

(...) To do this task, you need to have a Twitter account that meets the following requirements:- At least 100 followers - Your follower count needs to be at least double your following count (meaning - if you are following 100 users, you need to have at least $100 \times 2 = 200$ followers) - The majority of the 20 most recent tweets are in English - At least 8 out of 20 most recent tweets have no links, are Not Retweets, and sound natural and interesting.(...)

Campaign example 7: Twitter account quality rules

This is obviously requested in an effort to imitate a real Twitter user and to not seem like a newly created one. These tasks pay bonuses too, meaning that the pay can reach as high as \$ 0.25.

Posting (P) and Tweeting (T) were grouped together for being very similar. P and T is most prevalent on Twitter (716 campaigns; 80,211 tasks) and Google+ (485; 28,819). Connecting as a friend is mostly done on Facebook (13 campaigns; 9,190 tasks). Liking/Upvoting (L) is an activity performed on Reddit (539 campaigns; 77,104 tasks), Facebook (485; 71,200), Mixcloud and SoundCloud and Datpiff (143; 11,773), YouTube (63; 23,993), Instagram and Google+ (both 13 campaigns, 5,366 and 5,401 tasks, respectively), and on some other platforms (49).

The 10 biggest like/upvote campaigns ranged between 1,830 and 7,417 offered tasks (here, Facebook, Reddit, Instagram, Google+ campaigns were all in the top 10). The content of these top campaigns unfortunately were redacted using the rotator technique (see later in the article), but some of the remaining involved promoting persons not noted on Wikipedia (vanity-promoting); and some niche products. Size seems to be a limitation again, just like with search and engage campaigns: for celebrities with hundreds of thousands of followers, even as much as 7 thousand new likes hardly seem to matter, and the promotion technique does not seem to scale up to higher numbers.

This limitation might be not there in the case of comments (C) campaigns. The biggest comments campaign was conducted on YouTube, very similar to example 6, and it involved 21,140 tasks, three comments each, yielding more than 63,000 paid comments. The tenth-biggest involved 2,425 tasks, again three comments each. While a similar amount of likes would still represent just a fraction when it comes to comparing it to the likes received for the most popular YouTube videos, Facebook accounts, etc. For comments, the case is different, because only a small fraction of readers/visitors make comments. A cursory, non-representative investigation of the many YouTube videos reveals that it is very hard to find videos that have less than 20x viewers than comments. Thinking the other way around, 63,000 tentatious comments would suggest representing well over 1.2 million viewers, hence distorting the perception of actual viewers on what other's opinions are in relation to the topic.

However, in local communities with a smaller overall size, it seems that even small (L) campaigns can make a difference among the competition. Consider this example re-

lated to warriorforum.com (its self-description is: “The world’s largest Internet Marketing Community and Marketplace.”)

Title: Warriorforum Post: Comment Payment: 0.10

Number of workers accepted: 30

Job description:

Must have a Warriorforum account at least 2 months old, or a very active account if you joined recently.

1. Go to: www.warriorforum.com/<REDACTED>
2. Post a comment/testimony relating to the post **Must be original/no copying others post**

Tip on what you can post:

I learned a lot from the webinar,

I just join and I am excited, cannot wait to start working with <REDACTED>, this is an awesome program, etc.

Required proof that task was finished?

1. Name used to leave comment

Campaign example 8: Comments on a small forum

A cursory look at Warrior Forum reveals that the typical number of upvotes is around 10 and the number of comments (reply-s) is similar. The job offer above for 30 comments in example 8 would thus propel the entry among the top, while costing a mere \$ 3.3 for the client.

6.4 Smartphone apps

A distinct area of paid activity is installing smartphone applications, using/testing and rating them. There were 361 such campaigns with 16,934 tasks worth \$ 8,375.01. A typical campaign looks like this:

**Title: Android App Testing (<REDACTED>): Download+Install+Honest Review
Payment: 0.50**

Number of workers accepted: 370 Job description:

1. Install the app below: play.google.com/store/apps/<REDACTED>
2. Download the app
3. Open the app for 30 seconds and test
4. Rate the app

Optional: Leave 3, 4, or 5 stars

5. Write an honest review in the Microworkers proof box only

Required proof that task was finished?

1. Your Google username
2. Paste your review

Campaign example 9: Smartphone app ”testing”

We can interpret this campaign in several ways. The charitable interpretation would be that this is an honest test for the software. Even though the app is only required to run for 30 seconds, the freelancers would open it on dozens of different Android devices, with different capabilities, screens, API levels, etc. Thirty seconds is enough to run some self-assessment and report back to a server. This test could thus have some value from a Software Engineering perspective and may be a valid assignment.

However, this kind of test is surely already long overdue in the case of a published application. Issues around crashing apps and unwanted startup behavior should have been resolved way before then. The author speculates that these kinds of campaigns are instead promotional. The clients are usually careful not to explicitly order five stars or positive reviews (albeit there are counter-examples of such). Yet, in such a campaign, an initial, visible user base is created for the app. Again, we can assume that this is more useful in niches than in competition with mainstream applications, as the leading apps have millions of users already.

It must be mentioned that this kind of microtask carries security risks for the freelancer and for the general public. The fact that the freelancer installs apps for a fraction of a dollar presents an obvious opportunity for breaching their smartphones. Although the current number of tasks in these campaigns does not seem to be high enough, in theory it would be possible to create a zombie network for DOS attacks or similar purposes.

6.5 Signup

A very common type of campaign is the signup (U) campaign. These campaigns involve creating an account meeting some client-specified requirements. In many cases, the microworkers will be required to hand over the account credentials. Example 10 below was one of the biggest signup and account handover campaigns observed.

Title: YouTube: Create an Account

Payment: 0.10

Number of workers accepted: 2290

Job description:

YouTube: Create an Account

1. Go to www.youtube.com
2. Create a new account
3. Verify your email
4. Login to activate the account

Required proof that task was finished?

1. YouTube username or email
2. YouTube password

Your task will NOT be rated satisfied if your YouTube account requests phone verification.

Campaign example 10: Youtube Account Creation

The client in this case has acquired 2,290 YouTube accounts for a mere gross \$ 251.9. Example 2 from the introduction is a similar case, but for Gmail. The dangers posed by such campaigns are obvious. Besides promoting products, ideas and agendas, a cohort of 2,300 YouTube users can disrupt any smaller community on the platform and the use of fake accounts could facilitate the account owner to commit fraud or otherwise abuse the system.

What makes these mass account acquisitions very dangerous is that they are not easy to detect. Methods that are able to detect fake accounts typically only work if they are all created by the same person (Xiao u. a. 2015), but cannot be expected to work in this case as these accounts are created by real people. A landmark study by Gurajala u. a. (2015) involving the analysis of 62 million Twitter public user profiles relied on statistics about update frequency, reused profile pictures, and account creation days. Unfortunately, these factors can all be made to look genuine; for instance, the freelancers can be instructed to use profile pictures that are not reused; or possibly profile pictures themselves are acquired via microworkers (see example 13), and the creation times can be spread out with the help of “throttling”—a feature of the platform that allows only a certain number of tasks to be completed in a unit of time. Sometimes clients give instructions that enable the detection of such accounts, e.g., by requiring the freelancers to use the very same password. Also, it is probable that after handing over an account, the geolocation of the usage of that account is changed permanently, and so never again reflects the country of creation, which could be a factor in detection.

Not all signup campaigns seem to require account handover. For instance, the top four signup campaigns in terms of task numbers required a signup to two different website traffic providers and a polling site; involving 21,388, 19,952, 10,888, and 7,648 individual signups. Related to these campaigns was the biggest (99,999 workers) and cheapest (paying \$0.0) campaign observed, categorized as testing (Z), as technically it was a website spellcheck; its details are given in the following example:

Title: Qualification Test: Find the Misspelled Word

Payment: 0.0

Number of workers accepted: 99 999

This is a qualification for a future website test which will pay \$7.50 for about 11 minutes of work. This qualification test is to find workers with a great eye for detail.

1. Go to REDACTED URL
2. Find the misspelled word.
Hint: it is near the bottom Required proof that task was finished?
1. The wrongly spelled word in its wrongly spelled form

Campaign example 11: A recruitment campaign involving a small test to pass

All five sites (the aforementioned four traffic providers and the one in Example 11) were categorized as “Other” on the platform, and were not commonly featured. The scale of these campaigns explain why the category Other is so prominent in the platform aggregation in Table 4. Also, all the sites are basically recruiting microworkers for their own platform. The nature of tasks to be done there seems to be are traffic generation (visiting sites), participation in paid market research by answering surveys, etc.

Among the next five in the top 10 signup campaigns (places 6-10 ranked by the number of tasks on offer) was example 10, another account creation and handover involving 2,190 accounts to a site redacted with the rotator technique, plus three jobs requiring the signup of several thousands of users to various sites for unknown reasons.

6.6 Other interesting campaigns

This section covers several interesting campaigns that cannot easily be categorized in the other categories, many of them one-of-a-kind campaigns, and some of them seem rather strange and unexplained.

There was one campaign that requires the users to solve captchas. This is obviously to bypass a captcha-protected signup page. We can hypothesize that this is part of a human-in-the-loop automated account creator system.

There were several research campaigns observed. These are transparent and benign: the university or the research group is clearly present, there is usually a document attached as a brief for the research. The topic seems to be social psychology or web usability and ergonomics. The users are asked their gender and then made to do face expression recognition; evaluate risks; try out different webpage workflows, etc.

There is one observed snapchat promoter recruitment campaign (30 tasks x USD 0.5), see the following:

Objectives: I'm looking for cute girls who snap for marketing promotions (bonus possible).

Important: You must actually snap video and not just upload pictures from fake profiles.

1. Provide your Snapchat Username for proof (I will add you as a friend)

Campaign example 12: Snapchat recruitment

Some campaigns seem to be building stock photos, like example 13 below (1000 x \$ 0.11). Another project required photos of windshields. Yet another project asked for a selfie of the freelancer, and the consent to use it, but only from those who had no beard.

Important: You must agree to allow us to use your photo for promotional purposes in order to complete this task.

1. Take a well lit, clear photo of an office. **Important:** Make sure no people are in the photo.

Notes:- I need a clear photo of the office showing computers, desks, etc.- Photo should not be a photo from the Internet, we search for all photos on the Internet before we approve

Campaign example 13: Acquiring photos

Finally, for the following campaign there is just no explanation:

(30 x \$1.75) Write 12 lines of lyrics for an Anthem based on your own individual traditions and struggles (Make it relevant to your life today) (...)

Campaign example 14: An unexplained campaign

7 Techniques employed in campaigns

As explained in the Limitations section, there are a number of invite-only campaigns on the site, called “hire groups”. These allow a client to select the freelancers, as contrasted to public campaigns that are open to anyone to participate. Also, this allows a per-employee task customization by providing a spreadsheet of input variable values. While also being feature rich, hire group campaigns are usually hidden from the public view.

Rotators are another way of per-employee customization and also allows hiding the content of the campaign from public display.

Title: Forum: Sign up + Post + Screenshot Payment: 0.14

Number of workers accepted: 96

1. Go to this link: <bit.ly/<REDACTED URL ENDING1>>
2. Search for blogs from this search link
3. Find blog, website or forum you can post a comment on
4. Go to this link: <bit.ly/<REDACTED URL ENDING2>> and then copy comment from this page and post this comment in the website blog or forum

Required proof that task was finished?

1. Your Forum Username
2. URL of the comment
3. Screenshot of posting

Campaign example 15: The rotator technique

This technique allows the employer to customize the task per-employer without a hire group and to remove the instructions after the campaign is done without leaving a trace. Except for those freelancers who participated in the campaign, there is no way of knowing what sites, search keywords, or comments were involved in the job. The category “Redacted” among the platforms refer to this technique and not to data anonymization employed by the author in the examples in this article. Of course, there is no way of knowing if the client’s intention was just to rely on task customization, or to hide the campaign content, or both.

As explained at the Smartphone apps section, there might be ways for dressing up promotion campaigns as testing campaigns, by asking a couple of hundred users to install the app and then leave it there. Also, there might be search and engage campaigns masquerading as data collection and competition monitoring. In the case of some Amazon- and eBay-related campaigns, the freelancers are directed to search for different products, then to select from a given set of results, and then to collect prices, data, specifications, and to finally submit these as job proof. What makes these suspicious is that for an honest information campaign, it seems to be overly redundant to collect the same information many hundreds of times by many hundreds of microworkers. In reality, the point of these campaigns could be to make the microworker search and engage and then to spend time on the visited page while counting reviews and collecting information (the algorithm of a search engine might take the duration spent on a result page into account when adjusting itself), and then the accomplishment of the job can be conveniently verified by the client by looking at the collected data. Of course, these are just hypotheses for which there is no way to verify them.

Figure 1 summarizes the promotional methods observed, together with the supporting techniques featured in various gray promotional campaigns:

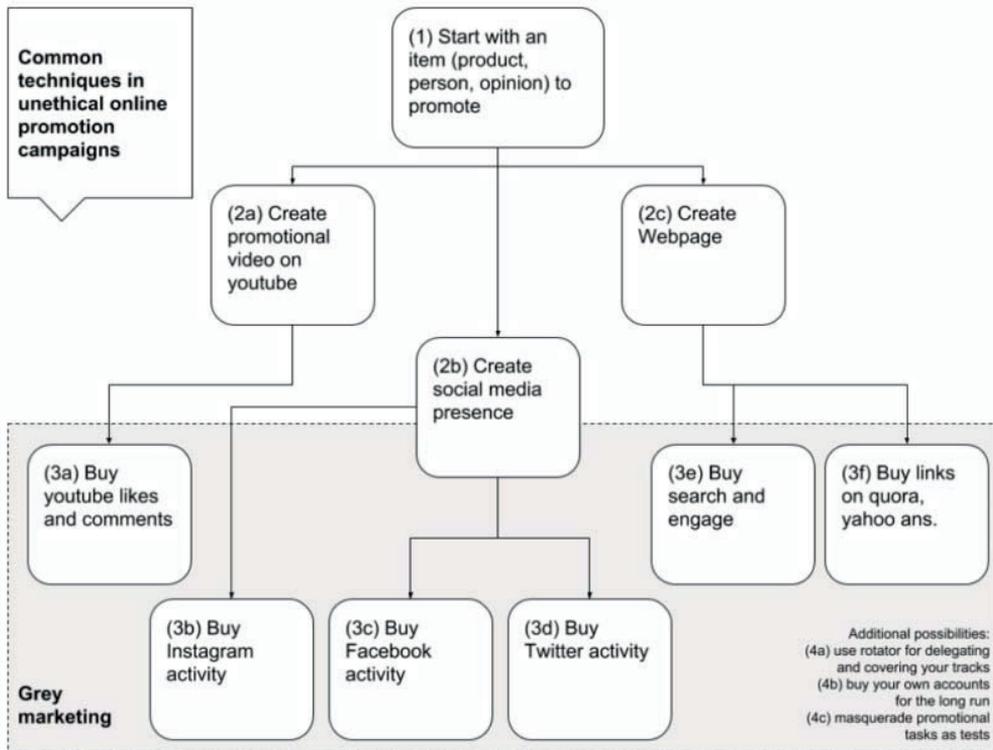


Figure 1: Elements of unethical online promotion campaigns.

8. Conclusion

This article provides insights into the black market of likes, upvotes, comments, retweets, votes in contests, and search engine manipulation. The subject of the investigation was microworkers.com, which is not a black market itself per se, but it has light regulation of its campaigns and so can be used by clients to participate in black-market activities. Also, it clearly is only one of several venues for running such campaigns. De Micheli und Stroppa (2013) have investigated several other players on the market (Fiverr, Seo-Clerks, Inter-Twitter, FanMeNow, LikedSocial, SocialPresence, SocializeUk, ViralMedia-Boost) and they have found that the market size is probably several millions of dollars, making the share microworkers.com a tiny fraction. Other sites even recruit on microworkers.com for similar microtasks. However, thanks to the fact that on microworkers.com the client has to orchestrate the campaigns itself, we can get an insight how the other players in the market, that sell complete like and follower packages, might be operating.

The nature of the microworkers.com campaigns was explained in the sections above. About their efficiency, we concluded that it probably varies. The main limitation is that it seems to be hard to purchase more than some tens of thousands of items. As explained in the section on Social media activity covering likes/upvotes (L), these numbers do not make a big difference when it comes to widely discussed political topics or celebrities, as in this

area millions of L items are not uncommon. However, in smaller communities, with normally dozens or hundreds of L items, they can make a huge difference. This is the context in which the effects of a total of 207,811 L tasks can be assessed. For instance, Reddit, on which 77,104 upvotes were purchased, is a platform where a couple of hundred or thousand purchased upvotes can go quite far, especially in thematic sub-Reddits. It might be noted that a similar number of downvotes would be much more significant as there are normally much less of these items—but no downvoting/dislike campaigns were observed.

For comments, we have to assume that the big observed campaigns, reaching 60,000 YouTube comments must be effective as these are quite high numbers when it comes to comments. It is of course unclear what the overall effect of tens of thousands of comments is on the thinking of the targeted audience. But it is enough to provide an apparent majority on almost any platform.

For online voting and contests, it seems that all kinds except the biggest contests can be rigged by microworker campaigns.

There are two areas where the efficiency is especially hard to assess: search and engage and app testing. For search and engage, over half a million tasks were observed and we must assume that the efficiency of these jobs really depends on the popularity of the topic in question. Also, the effect of these campaigns is really hard to track. In a similar way, a hypothesis was provided on how app developers on Android or iPhone might be trying to build an initial user base of a couple of hundred installs. The most popular applications have tens of millions of user and even their alternatives often have tens or hundreds of thousands (this also indicates just how hard the entry must be to that market). A better understanding of the app market places would be necessary to understand the significance of a couple of hundred individual users.

Finally, the knowledge that several thousand YouTube, Gmail, Snapchat and other accounts have been created and their usernames and passwords handed over during the period of observation is very troubling. Those accounts might be effectively used for large scale gray promotion campaigns and could also pose a security threat at the same time.

Future work could involve participatory research as a freelancer on this or other platforms to reveal the experience of a microworker as well as to discover more about the invite-only/hire-only campaigns. In cases of comment copy-pasting, the source and nature of the comments could also be learned. Another area could be investigating the logic and goals behind the traffic generator sites and unions that similarly to microworkers' sites rely on the completion of menial tasks.

Acknowledgement

This work was supported by the MTA Bolyai scholarship and the ÚNKP-19-4 New National Excellence Program of the Ministry for Innovation and Technology.

References

- Allcott, H. and M. Gentzkow, Social media and fake news in the 2016 election, *Journal of Economic Perspectives* 31. (2017) 2., 211–36.
- Arthur, C., How low-paid workers at 'click farms' create appearance of online popularity, *The Guardian* . 2 August 2013.
<https://www.theguardian.com/technology/2013/aug/02/click-farms-appearance-online-popularity>
- Buhrmester, M., T. Kwang and S. D. Gosling, Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data?, *Perspectives on psychological science* 6 (2011) 1., 3–5.
- Buttcher, S., C. L. Clarke and G. V. Cormack, *Information retrieval: Implementing and evaluating search engines*, Mit Press, 2016.
- Caldwell, C. , Not being there, *The New York Times*. 12 August 2007.
- Clark, J. , Google turning its lucrative web search over to AI machines, *Bloomberg Technology* 26. (2015).
- Cook, D. M., B. Waugh, M. Abdipannah, O. Hashemi and S. A. Rahman, Twitter deception and influence: Issues of identity, slacktivism, and puppetry, *Journal of Information Warfare* 13. (2014) 1.; 58–IV.
- Crone, D. L. and Williams, L. A. (2017). Crowdsourcing participants for psychological research in australia: A test of microworkers, *Australian Journal of Psychology* 69(1): 39–47.
- De Micheli, C. and A. Stroppa, Twitter and the underground market, *11th Nexa Lunch Seminar*, Vol. 22. (2013), Politecnico di Torino.
- Del Riego, A. , Digest comment-context for the net: A defense of the ftc's new blogging guidelines, *JOLT Digest, an online companion to the Harvard Journal of Law and Technology*. 2009.
- Forrest, E. and Y. Cao, Opinions, recommendations and endorsements: The new regulatory framework for social media, *Journal of Business and Policy Research* 5. (2010), 2. 88–99.
- Gardlo, B., M. Ries, T. Hobfeld and R. Schatz, Microworkers vs. facebook: The impact of crowdsourcing platform choice on experimental results, *Quality of Multimedia Experience (QoMEX), 2012 Fourth International Workshop on*, QoMEX, 2012, pp. 35–36.
- Gurajala, S., J. S. White, B. Hudson and J. N. Matthews, Fake twitter accounts: profile characteristics obtained using an activity-based pattern detection approach, *Proceedings of the 2015 International Conference on Social Media & Society*, Social Media Society, 2015, p. 9.
- Hirth, M., T. Hobfeld and P. Tran-Gia, Anatomy of a crowdsourcing platform using the example of microworkers. com, *Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS), 2011 Fifth International Conference on*, IMIS, 2011, pp. 322329.
- Howe, J. The rise of crowdsourcing, *Wired magazine*, 14. (2006).
- Laperdrix, P., W. Rudametkin and B. Baudry, Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints, *Security and Privacy (SP), 2016 IEEE Symposium on*, IEEE, 2016, pp. 878–894.
- Nguyen, N. , Microworkers crowdsourcing approach, challenges and solutions, *Proceedings of the 2014 International ACM Workshop on Crowdsourcing for Multimedia*, ACM, 2014, pp. 1–1.
- Nicholas Confessore, R. H. Gabriel J.X. Dance and M. Hansen, The follower factory, *The New York Times* . 31 January 2018.
- Paolacci, G., J. Chandler and P. G. Ipeirotis, Running experiments on amazon mechanical turk, *Judgment and Decision Making* 5. (2010) 5.
- Rutz, O. J. and R. E. Bucklin, From generic to branded: A model of spillover in paid search advertising, *Journal of Marketing Research* 48. (2011), 1., 87–102.
- Smith, D. (2018). Putin's chef, a troll farm and russia's plot to hijack us democracy, *The Guardian* . 17 February.
<https://www.theguardian.com/us-news/2018/feb/17/putins-chef-a-troll-farm-and-russias-plot-to-hijack-us-democracy>

-
- Thackeray, R., B. L. Neiger, C. L. Hanson and J. F. McKenzie, Enhancing promotional strategies within social marketing programs: use of web 2.0 social media, *Health promotion practice* 9. (2008) 4., 338–343.
- Xiao, C., D. M. Freeman and T. Hwa, Detecting clusters of fake accounts in online social networks, *Proceedings of the 8th ACM Workshop on Artificial Intelligence and Security*, ACM, 2015, pp. 91–101.
- Yang, S. and A. Ghose, Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence?, *Marketing Science* 29. 2010, 4., 602–623.
- Zhang, Y., X. Li and T.-W. Wang, Identifying influencers in online social networks: The role of tie strength, *International Journal of Intelligent Information Technologies (IJIT)* 9. 2013, 1., 1–20.

