Bad Job: Abusive Work on Alternative Microtask Platforms

Christina Yeung, Daniel Salsburg, Joseph A. Calandrino
Office of Technology Research and Investigation
Federal Trade Commission

***Abstract*—Microtask platforms aim to pair employers and workers to complete small tasks for modest pay, but the purposes of these tasks are not always benign. Researchers have identified problematic tasks on popular microtask platforms like Amazon Mechanical Turk, and such tasks may be relatively more common on smaller platforms. Recent work examining these smaller alternative platforms is limited, and the nature of the work on these platforms evolves with the desires of employers and the practices of the platforms. We provide an up-to-date view of the work available via alternative microtask platforms. To do so, we collected details from three alternative platforms over approximately a month, categorizing the available work. We find that potentially abusive work persists in well-known categories like search engine optimization, but we also uncovered new and emerging categories of work, such as tasks that may manipulate spam filters. We comprehensively explore these categories and discuss potential mitigation approaches.**

I. Introduction

Microtask platforms connect employers with workers to complete often short, repetitive tasks online for modest pay. Tasks range from participating in surveys to classifying images. Unfortunately, these tasks are sometimes abusive [24], [35]: a worker's efforts might go toward generating fake product reviews or undermining security measures like CAPTCHAs. Although abuse of larger platforms remains a concern, some questionable work may be migrating towards smaller alternative platforms [12], [48]. Recent analysis of the tasks available to workers on these smaller platforms is limited. We provide an updated view on the nature of work on these platforms, exposing new details and categories of potentially abusive work.

On microtask platforms, employers request work by creating *campaigns*, which describe a task to be completed one or more times. For example, an employer might wish for 50 people to complete a survey. The employer would post a campaign for the survey, listing 50 available tasks. Via the platform, employers provide details of a campaign, including title, potential earnings, expected task completion time, number of requested tasks, and any additional instructions. Workers review campaigns and complete tasks, receiving payment upon completion. Payment occurs through the platform with logistics and fees varying by platform. Platforms typically collect fees at minimum from employers, who may incur a nominal fee when creating a campaign or paying a worker.

Past research has identified campaigns driving spam, fraud, and other problematic activity on popular microtask platforms like Amazon Mechanical Turk [24], [35] (see Section II). The use of distributed workers may undermine basic measures—from bot detection to coarse rate limits—seeking to thwart this activity. Questionable work may be a relatively greater issue on smaller platforms [48]. While some older studies have explored those platforms [12], [22], neither the platforms nor those seeking to exploit them are static. As platforms' practices, online threats, and other factors evolve, workers may encounter a changing landscape of available work.

We explore the current state of alternative microtask platforms. We collected details of campaigns on three platforms' websites for approximately a month (see Section III). Through a mix of manual and automated analysis, we categorized the 2,609 observed campaigns by the nature of the requested work and ultimately manually verified each category.

We find that a vibrant, active ecosystem for abuse endures on alternative platforms (see Section IV). Campaigns seemingly aimed at manipulating search results, social media, and online reviews are common on these platforms. We analyze categories of work, covering forms that prior research had not identified on alternative platforms. For example, we observed campaigns that involve marking emails as not spam or indicating that certain ads are helpful. We also observed campaigns that may abuse workers themselves.

The abundance of potentially abusive campaigns raises questions about the adequacy of measures that alternative platforms employ to prevent such work. We discuss the possible impact of these campaigns and mitigation strategies for platforms, workers, and affected third parties (see Section V). We also explore areas in which future research may help us to understand and address abusive work (see Section VI).

II. Related Work

Beyond exploring abusive campaigns on microtask platforms, prior work has studied related abuse more broadly and examined additional details of microtask work.

Abuse on Microtask Platforms. Choi et al. [12] examined five platforms of varying sizes and developed a system to classify campaigns as malicious or legitimate. In 2011, Motoyama et al. [35] analyzed Freelancer activity, finding that 31% of campaigns related to spam, search engine optimization (SEO), or artificial social media activity. The same year, Wang et al. [48] studied a variety of microtask platforms. While they noted a sharp decrease in observed abusive campaigns on Amazon Mechanical Turk between 2010 (41%) and 2011

(12%), they estimated that 70%-95% of campaigns were malicious on four US-based alternative platforms. Given this and the intervening decade, we provide an updated view focused exclusively on alternative platforms. We find that potential abuse persists and identify novel forms.

Other past work focuses on specific categories of abuse via microtask platforms, including manipulation of reviews [16], [31], [38], social media [29], [30], [32], [42], and shopping search results [45]. Prior work also explores worker privacy on Mechanical Turk, including concerns [49] and protective behavior [40], [41].

Related Abuse. Malicious microtasks often drive known forms of abuse. Prior work explores fake reviews [26], search engine manipulation [4], [13], [14], fake social media activity [20], survey scams [28], and click fraud [50], including the actors' goals and economics [5], [44]. Existing work also proposes and examines techniques to detect, prevent, or mitigate abuse, including fake account creation, social media spam, and fake reviews [7], [10], [11], [17], [18], [36], [47].

Microtask Work. Prior work studies microtasks—especially on Mechanical Turk—from various other perspectives, including worker demographics [23], [39], estimated pay [8], [21], and crowdsourced data quality [6], [9], [19], [43]. Irani et al. [25] explore the ethics of microtask work and develop a tool to increase transparency. Hirth et al. [22] consider alternative platforms but focus on users and use of the platforms.

III. Approach

The goal of this study was to explore the ecosystem of alternative microtask platforms, examining the nature of available work. To do so, we monitored open campaigns on three platforms for 30 days. We categorized campaigns by requested work, manually refining all assignments.

A. Platform Identification

Mimicking a worker, we began with web searches for crowdsourcing and microtask platforms. We checked results, including lists and reviews. Unlike broader gig and freelance work, relevant platforms offer small, discrete tasks that workers complete online. We considered English-language sites for which employers post campaigns and set pay for task completion. To focus on alternative platforms, we omitted sites in the Alexa top 5,000.

This yielded seven platforms. During the study period, four platforms' sites used or introduced measures like CAPTCHAs to prevent crawling.[1] We respected these measures by excluding these sites, but this may have introduced bias. The remaining sites were JobBoy, Microworkers, and Minijobz.[2]

These sites' WHOIS registration dates are from 2007 to 2010. On February 20, 2020, JobBoy claimed more than 300,000 workers [27], and Microworkers claimed over 1.5 million workers and 45 million completed tasks [34]. On March 16, 2020, Alexa rankings of JobBoy, Microworkers, and Minijobz were 269,621, 9,195, and 572,741 [3].

These platforms provide similar site structures and campaign details. Each supplies a list of open campaigns with metadata about each campaign, including title, pay, expected task completion time, and the remaining number of tasks available. Each campaign has a unique ID and a webpage on the platform. Beyond metadata, this page offers optional descriptive text for instructions and other details. Employers can increase the requested tasks for open campaigns.

B. Campaign Crawler

We built a Python-based crawler using Selenium WebDriver and Requests.[3] We adapted it to each platform's site. Our crawler visits a site, paginates through the campaign list, and extracts details for each campaign. The crawler then visits each campaign's webpage, obtains descriptive text, and stores HTML and a screen capture. To avoid overburdening platforms, we implemented rate limiting such that we visited no more than four pages per platform each minute. Our tool identifies URLs in descriptive text, attempts to visit them, and records redirects. If the final page loads successfully within 30 seconds, the crawler saves the HAR file,[4] a screenshot, and HTML.

We ran the crawler daily from February 19, 2020 to March 19, 2020 inclusive. Due to a network issue, collection failed March 1, 2020. We collected details of campaigns each day they were active. We exclude cases in which a campaign webpage request yielded a persistent server error, an employer suspended a campaign, or the platform listed the campaign as complete. In general, campaigns disappear from platforms when tasks are complete, but we cannot know why a campaign disappears.

C. Campaign Categorization

Using unsupervised learning, we classified a single day's campaigns. We reviewed the results to identify keywords for meaningful categories. We then tagged all campaigns with the keywords before manually refining assignments.

Preliminary Classification. We applied unsupervised learning to campaigns from a single randomly chosen day. Using NLTK,[5] we lemmatized the text of each campaign's title and description. We sought to resolve URLs, following redirects. If a URL failed to resolve, we excluded it. Otherwise, we replaced it with the resolved URL's domain, substituting spaces for periods (e.g., `www.example.com/path` would become "example com"). We then applied TF-IDF, clustering the results using k-means. Based on previously identified categories [35], [48], we set k to 25. This yielded a preliminary classification of the campaigns, with each in a single class.

[1] Clickworker (clickworker.com), Online Micro Jobs (onlinemicrojobs.com), Picoworkers (picoworkers.com), and Rapidworkers (rapidworkers.com)

[2] https://jobboy.com, https://microworkers.com, https://minijobz.com

[3] https://www.selenium.dev/, https://requests.readthedocs.io/

[4] https://w3c.github.io/web-performance/specs/HAR/Overview.html

[5] https://www.nltk.org/

TABLE I
PER-PLATFORM OBSERVATIONS DURING STUDY PERIOD.

	Campaigns		Tasks			Pay		Time (min)		Crawls
Platform	Total	New	Mean	New	Completed	Mean	Newly Paid	Mean	Newly Spent	Mean
JobBoy	202	1	43.8	143	768	\$0.22	\$118.44	4.0	3,189	25.3
Microworkers	2,382	1,189	250.2	267,761	92,479	\$0.26	\$13,067.24	7.7	751,509	12.5
Minijobz	25	0	14.7	0	7	\$0.39	\$4.15	10.2	30	25.6
Overall	2,609	1,190	232.0	267,904	93,254	\$0.26	\$13,189.83	7.4	754,728	13.7

Manual Review. Considering all stored data for the campaigns, we manually analyzed the preliminary classification. We merged and split buckets to yield meaningful groups, and we identified keyword combinations that tended to be unique to groups, such as *search* and *result* for an SEO-related group.

Labeling and Refinement. We used the keywords to assign initial categories to all campaigns. Campaigns can be in multiple categories, and some categories are subcategories, such as different forms of social media activity. For campaigns in subcategories, we labeled the campaigns with both the subcategory and the overarching category.

Our goal was an accurate view of campaigns. Therefore, we used all collected data to review and refine categories for each campaign manually. For example, we identified a new category by analyzing URLs and redirects. Manual review ultimately determined campaigns' categories. This resulted in 18 categories, including subcategories (see Section IV-A).

Appendix A provides our keywords. While useful for our purpose—streamlining manual refinement—we caution against assuming that they provide a robust or generalizable classifier. Our goal was not to produce a classifier, so we did not separate training and test data or perform cross-validation.

IV. FINDINGS

We observed 2,609 campaigns. Table I provides platform-specific details on the number of observed campaigns as well as per-campaign means for the number of observed tasks, advertised pay per task, employer-estimated time per task, and number of crawls in which we observed the campaign. It also offers details on observed activity, including new campaigns (i.e., campaigns created during the observation period), tasks added/completed, estimated total payments to workers (before fees), and estimated time spent on completed tasks. Many numbers are likely underestimates. Campaigns typically disappear when tasks are complete. We cannot confirm completed tasks from a campaign's final day, and campaigns might come and go between crawls.

Microworkers hosted an order of magnitude more campaigns than JobBoy and nearly two orders more than Minijobz. Open campaigns decreased from 1,419 to 762 over our observation period, largely due to a single employer. Appendix B discusses the long-term trend in open campaigns, and Section IV-A explores the outsized impact of some employers.

A. Distribution of Categories

Table II provides the distribution of categories over campaigns. Like Motoyama et al. [35], we find the SEO category to be most common, with 1,526 campaigns. Social media is the second most common. Appendix C offers per-platform details.

Microworkers and Minijobz provide an employer user ID with campaigns. One employer posted over half of all SEO campaigns (875/1,526) and over 40% of all Microworkers campaigns (1,033/2,382). Many of these campaigns have only minor variations—such as similar instructions with different search terms—and they typically promote an online check-writing service. On Minijobz, three employers were responsible for more than a quarter of observed campaigns (7/25).

B. Category Analysis

For each category, Table II also provides the per-campaign mean number of observed tasks, advertised pay per task, estimated time per task, and number of crawls in which we observed the campaign. We explore each category here.

SEO. SEO seeks to influence search results, promoting webpages, products, or other content. Accepted means of promoting content exist, but observed campaigns in the three subcategories appear to drive largely inauthentic behavior:

- *Search term.* These campaigns instruct workers to enter given search terms and visit particular results. 77% direct workers to search via Google services, excluding YouTube. Other search destinations include YouTube (11%) and retail sites (3%), like Amazon and eBay.
- *Timed activity.* These campaigns request activity like visiting a search result or keeping software installed for a given period. This may increase the likelihood that the workers' interest appears legitimate [33]. 75% of observed campaigns specify a length of time, from 2 seconds (visiting a search result) to over 4 weeks (keeping a web browser installed). We are unaware of past research identifying this category on alternative platforms.
- *Backlink.* These campaigns provide a URL for workers to post elsewhere. These links may influence the content's position in search results [37]. This category is less common but higher paying than other SEO subcategories.

Social media. The 1,300 campaigns in these two subcategories request activity on online social networks:

- *Engagement.* These campaigns have workers promote content on social media, such as sharing or "liking"

TABLE II
DISTRIBUTION OF CAMPAIGN CATEGORIES AND METADATA (MEANS) FOR CAMPAIGNS IN EACH CATEGORY.

Category	Campaigns	Description	Tasks	Pay	Time (min)	Crawls
SEO (All)	1,526	Influence search results	245.1	\$0.19	7.4	14.0
—Search term	998	Enter search term, interact with a result	280.4	\$0.20	6.5	14.2
—Timed activity	859	Perform activity for a minimum duration	203.4	\$0.19	7.2	12.9
—Backlink	36	Post links to content	171.3	\$0.48	10.6	11.3
Social media (All)	1,300	Increase visibility/reach on social media	245.1	\$0.16	8.2	13.5
—Engagement	1,275	Like, share, retweet, etc. content	241.3	\$0.16	8.3	13.6
—Connections	686	Add friends, followers, subscribers, etc.	191.0	\$0.13	8.9	13.3
Personal information	422	Provide personal information	167.5	\$0.44	5.0	13.9
Review activity	231	Write or interact with reviews	187.6	\$0.46	5.8	9.2
Sign up	219	Sign up for account, offer, etc.	116.0	\$0.56	5.5	18.6
Account use	181	Use account to complete task	285.8	\$0.33	6.4	8.2
Redirect	136	Visit link resolving to different pages	42.1	\$0.38	4.5	21.6
Download/install	122	Download/install apps, software, files, etc.	83.2	\$0.90	7.3	17.4
Shopping	69	Add items to cart, wish list, etc.	478.6	\$0.30	5.7	9.7
Spam filter	65	Mark email as “not spam”	271.5	\$0.26	6.3	4.2
Gift card	60	Complete task for gift card offer	23.5	\$0.44	4.6	20.2
Account purchase	50	Create and provide accounts	270.0	\$0.80	11.9	10.1
Ad activity	26	Interact with ads	73.6	\$0.40	5.2	7.5
Uncategorized	110	—	244.0	\$0.31	21.1	8.2

content. This activity might increase content’s reach, reputability, and algorithmic placement. 76% do not mandate a social network, but the most observed social networks are Facebook (26%), Twitter (19%), and YouTube (13%). Less than 5% set requirements for the worker’s social media account, like a minimum number of friends.

- *Connections.* These campaigns generate social media connections, such as “friend” or follower connections. These connections may improve an account’s reputation, influence, and reach. 96% of these campaigns are also in the engagement category. These connections may be less valuable than organic followers [35] but may evade detection by platforms [32]. The fake follower economy on social networks is well studied [20], [44].

77% of social media campaigns are also SEO campaigns, having workers take steps like searching for and sharing items.

Personal information. These campaigns instruct workers to submit personal information, such as email address (93%), name (25%), and ZIP code (7%). Some campaigns also require workers to submit a photo ID (3%) or complete “Know Your Customer” steps (1%), and one requires bank account details.

The purpose of this collection may vary from limited campaign-specific needs to lead generation or even identity theft. 29 campaigns appear to provide referral codes—including all campaigns requesting “Know Your Customer” and bank account details—and 16 involve Robinhood. If referral code issuers like Robinhood did not authorize these campaigns, employers may be abusing referral programs for commissions. In cases like referral abuse, an employer might not obtain the worker’s information, but the worker nevertheless must provide the information to complete the task.

Review activity. These campaigns relate to ratings or reviews. 79% ask workers to write reviews, but campaigns also include requests to rate items (18%), “upvote” comments (9%), indicate that ads on social media are helpful (5%), and perform other activity. This might affect customers’ perceptions of items and items’ relative visibility.

Top destinations are social media platforms (57%, including YouTube) and app stores (15%), but others include Amazon (5%) and review platforms (e.g., TripAdvisor; 3%). 183 campaigns request written reviews, but only 6 explicitly instruct workers to indicate sponsorship. Many campaigns (61%) request honest reviews or make neutral requests, but even some of these offer mixed messages: 4% of campaigns separately request both honest and positive feedback. Some campaigns (11%) ask workers to share negative feedback privately with the employer, which may effectively bury negative reviews.

Sign up. These campaigns instruct workers to sign up for services, accounts, and offers (e.g., sweepstakes). Because the sign-up process often requires details like contact information, 62% also fall in the personal information category.

These campaigns sometimes direct workers to other income opportunities, like other gig economy platforms and paid survey sites. 14% lead workers to a social media site as an intermediate step, such as clicking a link in a social media post. 8% send workers to sites in which they could sign up to receive gift cards or to be entered in sweepstakes.

Account use. These campaigns require workers to use an account to complete a task. A worker may need to create an account, but the ability to use an existing account suggests that creation is not the primary goal. 52% request that workers access a Google account, specifying Gmail (30%), YouTube (18%), or simply Google (5%). Retail and social media sites are also popular, including Amazon (8%) and Facebook (5%).

Account use campaigns usually overlap with other categories (99%). Campaigns’ goals depend on those other categories and vary from social media activity to shopping activity.

Redirects. These campaigns provide URLs that can lead to different pages. One worker might see a sign-up page for Hulu;

another might see a sweepstakes offer. Instructions are often inaccurate. We suspect these campaigns posed challenges for workers, including difficulty proving task completion.

We considered resolved URLs each day a campaign was available. If a URL resolved to different pages, we examined the series of redirects, labeling the source of the split a *redirect service*. Campaigns in this category have URLs leading to identified redirect services. 68% rely on three services (see Appendix D) that describe themselves as "cost-per-action networks." They show third-party ads and receive payment if a visitor takes a desired action, like signing up for a service [1]. These campaigns persisted over more daily crawls than average, suggesting they remain active unusually long.

Download/install. These campaigns instruct workers to download and sometimes install items from mobile apps to eBooks. Top destinations are Apple's App Store (32%) and Google Play (18%). 7% of campaigns instruct workers to download web browsers, typically a stock web browser such as Brave. Based on categories provided by destinations (e.g., app stores), observed campaigns routinely require download of finance-related items ($\geq$ 16%) and games ($\geq$ 12%). Campaigns in this category had the highest mean pay ($0.90), perhaps to overcome concerns about downloading items.

The goals of these tasks may vary from inflating the popularity of an item to exploiting affiliate programs. For example, we observed tasks that require users to download a web browser and keep it installed for a period of time, providing an affiliate code in a link to the browser.

Shopping. These campaigns have workers make purchases (4%) or add items to shopping carts (84%), wish lists (7%), or watch lists (4%). Many direct workers to Amazon (84%), but destinations are diverse (see Appendix E). Campaigns routinely involve multiple items, such as adding multiple products to a cart. Workers select products through customized links (48%), search instructions (43%), or social media posts with product links (9%). Su et al. [45] explored similar campaigns.

Spam filter. These campaigns instruct workers to mark emails as "not spam," possibly undermining spam filters. 77% of observed campaigns require Gmail, and 6% encourage it.

97% provide a form page for workers to enter an email address. After submission, the form gives instructions, such as email subject lines. 3% direct workers to sign up for discount or deals lists, and workers must ensure that relevant emails appear in their inboxes. Workers may need to monitor for multiple emails. These campaigns seem to be short-lived: on average, we observed them over relatively few of our daily crawls (4.2). We are unaware of past research identifying this category of work on alternative microtask platforms.

Gift cards. These campaigns offer gift cards for completing tasks. Specified gift card values often greatly exceed task pay. 65% lead to websites that present gift card and other offers for completing surveys. Sites and survey questions may drive traffic to affiliate sites or collect personal details. Once a survey is complete, workers routinely see product offerings and new surveys to finish before they can claim rewards. Aspects of these campaigns resemble survey scams that Kharraz et al. [28] explored. 73% are also redirect campaigns.

Based on appearances in crawls, these campaigns remain active for relatively long. Oddly, they had the smallest mean number of tasks, above-average pay, short employer-estimated completion time, and promises of gift cards. Thus, workers should want to flock to these quick, high-paying tasks. Two possible explanations for this not occurring are that workers avoided these campaigns or could not prove task completion.

Account purchase. For these relatively high-paying campaigns, workers create online accounts and provide them to the employer. This may subvert defenses against bulk account creation. 86% request email accounts, and 74% specify Gmail. 18% provide a password, 46% specify a recovery account, and 44% simply ask workers to share account credentials.

Ad activity. These campaigns require workers to interact with ads, including clicking on ads (54%) and specifying whether ads on social media are helpful (46%). Depending on the specifics, these campaigns risk requesting click fraud.

Uncategorized. Under 5% of campaigns remained uncategorized. Some provide little more than links to external pages with minimal information. Some are unique, like an apparently legitimate academic survey. These campaigns have a much longer mean estimated completion time than any category.

V. Discussion

Many observed campaigns appear manipulative. Manipulation could harm users—who see less useful search results, reviews, and other content—and businesses seeking to compete fairly. Emerging types of campaigns appear to abuse user feedback for spam filters or to generate artificial interactions with ads on social media, but abuse could target any system relying on user activity.

Campaigns also may abuse workers. Like those on larger platforms, campaigns on alternative platforms may collect worker data. Some campaigns also involve a seemingly endless series of surveys or unclear tasks. We are uncertain how workers could prove completion of many such tasks.

A variety of techniques can help to mitigate abusive campaigns on alternative microtask platforms. Abusive campaigns may be just one part of a larger operation—such as a spam operation—which may offer additional points of intervention.

Platforms. Microtask platforms manage participants and campaigns as well as the interface and interactions mediating work. They also have substantial insight into activity on the platform. To help workers identify problems, platforms may wish to provide employer details like user ID, past campaigns, and payments [46]. Platforms could also encourage meaningful campaign instructions and details of any data collection and use. Platforms should be mindful of abuse and incentives. An employer may seek to undermine metrics, and a worker may lose pay by reporting an issue mid-task [40].

A single Microworkers employer posted more than 40% of campaigns. For major employers, platforms could periodically inspect campaigns, payment patterns, and complaints. Similarly, some campaigns fall in clusters that may be worth analyzing: many observed SEO campaigns differ primarily in search terms. Platforms might also wish to monitor for suspicious account creation, task completion, and other patterns.

If a platform knowingly assists employers seeking to break the law, its operators may wish to consider the possible legal implications of that choice [15].

Other Companies. Companies should factor abuse from these platforms into their threat models for both new and existing systems. Li et al. [31] note evidence of fake reviews from abusive campaigns. Similar evidence—such as bursts of suspicious activity—may exist for other forms of abuse. Companies could also monitor alternative platforms proactively. While small, these platforms can generate considerable activity. In one month, we observed more than 90,000 completed tasks.

Malicious campaigns may drive known forms of abuse, such as affiliate fraud. Improved countermeasures against these forms of abuse—ideally accounting for distributed workers—may also help protect against abusive campaigns.

Workers. When considering tasks, a prospective worker may rely heavily on employer-provided campaign details. Unfortunately, problematic campaigns might not be obvious from those details alone. Worker communities have emerged around the Mechanical Turk platform. These allow workers to share experiences and guidance [2]. Workers on alternative platforms may benefit from embracing similar forums.

VI. Conclusion

Our findings suggest that potentially abusive work remains abundant on alternative microtask platforms. Observed campaigns may drive inauthentic activity towards new goals—such as manipulating spam filters—or harm workers themselves.

Additional details or patterns might become apparent if considering other platforms or collecting data over different intervals and periods. For example, Microworkers was by far the most active platform we examined, but three platforms we excluded due to anti-crawling measures also had Alexa rankings below 20,000 on March 9, 2021 [3]. These similarly popular alternative platforms might offer a complementary perspective. Future researchers might wish to consider means of understanding any abusive campaigns on these platforms while respecting legitimate goals of anti-crawling measures.

Studies examining the relationships between parties in this ecosystem could help us understand the role that alternative platforms, employers, and other parties may play in online abuse more generally. In some cases, it may be feasible to connect campaigns on alternative platforms with activity elsewhere, such as linking accounts sold on underground forums to accounts created via microtasks.

Exploration of workers and their experiences on alternative platforms could help us better understand worker recruitment, possible harms to workers themselves (e.g., non-payment or identity theft), and interventions that might mitigate abuse. We would also benefit from broader exploration of mitigation techniques against abusive campaigns, including the efficacy of existing and potential approaches by different parties.

Acknowledgments

We thank Phoebe Rouge for assistance developing our web crawler and Tina Del Beccaro, Dana Roman, and Phillip Miyo for assistance setting up our collection infrastructure. We also thank Laura DeMartino, Hugh Huettner, and our anonymous reviewers for their helpful feedback.

References

[1] "Cost per action (CPA)," *MarketingTerms.com*, Accessed: Oct. 16, 2020, https://www.marketingterms.com/dictionary/cost_per_action.

[2] "Turker Nation: Helping mTurk workers make more money on Amazon Mechanical Turk," *Reddit*, Accessed: Oct. 16, 2020, https://www.reddit.com/r/TurkerNation/.

[3] Alexa, "Alexa - competitive analysis, marketing mix, and website traffic," https://www.alexa.com/siteinfo, last accessed: Mar. 9, 2021.

[4] A. Allam, P. Schulz, and K. Nakamoto, "The impact of search engine selection and sorting criteria on vaccination beliefs and attitudes: Two experiments manipulating Google output," *Journal of Medical Internet Research*, vol. 16, p. e100, 04 2014.

[5] M. Anderson and J. Magruder, "Learning from the crowd: Regression discontinuity estimates of the effects of an online review database," *The Economic Journal*, vol. 122, 09 2012.

[6] T. Behrend, D. Sharek, A. Meade, and E. Wiebe, "The viability of crowdsourcing for survey research," *Behavior Research Methods*, vol. 43, pp. 800–13, 03 2011.

[7] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida, "Detecting spammers on Twitter," *Collaboration, Electronic messaging, Anti-Abuse and Spam Conference (CEAS)*, vol. 6, 01 2010.

[8] J. Berg, M. Furrer, E. Harmon, U. Rani, and M. S. Silberman, *Digital labour platforms and the future of work: Towards decent work in the online world.* ILO Publications, 09 2018.

[9] A. J. Berinsky, G. A. Huber, and G. S. Lenz, "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk," *Political Analysis*, vol. 20, no. 3, p. 351368, 2012.

[10] Q. Cao, M. Sirivianos, X. Yang, and T. Pregueiro, "Aiding the detection of fake accounts in large scale social online services," in *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI '12)*, 2012.

[11] Q. Cao, X. Yang, J. Yu, and C. Palow, "Uncovering large groups of active malicious accounts in online social networks," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*, 2014.

[12] H. Choi, K. Lee, and S. Webb, "Detecting malicious campaigns in crowdsourcing platforms," in *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016, pp. 197–202.

[13] R. Epstein and R. E. Robertson, "The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections," *Proceedings of the National Academy of Sciences*, vol. 112, no. 33, pp. E4512–E4521, 2015. [Online]. Available: https://www.pnas.org/content/112/33/E4512

[14] R. Epstein, R. E. Robertson, D. Lazer, and C. Wilson, "Suppressing the search engine manipulation effect (SEME)," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, pp. 1–22, 12 2017.

[15] L. Fair, "VoIP where prohibited by law?" *Federal Trade Commission*, Jan. 30 2020 (accessed Oct. 16, 2020), https://www.ftc.gov/news-events/blogs/business-blog/2019/12/voip-where-prohibited-law.

[16] A. Fayazi, K. Lee, J. Caverlee, and A. Squicciarini, "Uncovering crowdsourced manipulation of online reviews," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15)*, 2015.

[17] H. Gao, J. Hu, C. Wilson, Z. Li, Y. Chen, and B. Y. Zhao, "Detecting and characterizing social spam campaigns," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC '10)*, 2010.

[18] S. Ghosh, B. Viswanath, F. Kooti, N. K. Sharma, G. Korlam, F. Benevenuto, N. Ganguly, and K. P. Gummadi, "Understanding and combatting link farming in the Twitter social network," in *Proceedings of the 21st International Conference on World Wide Web (WWW '12)*, 2012.

[19] J. Goodman, C. Cryder, and A. Cheema, "Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples," *Journal of Behavioral Decision Making*, vol. 26, pp. 213–224, 07 2013.

[20] C. Grier, K. Thomas, V. Paxson, and M. Zhang, "spam: The underground on 140 characters or less," in *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS '10)*, 2010.

[21] K. Hara, A. Adams, K. Milland, S. Savage, C. Callison-Burch, and J. P. Bigham, "A data-driven analysis of workers' earnings on Amazon Mechanical Turk," in *Proceedings of the 2018 ACM CHI Conference on Human Factors in Computing Systems (CHI '18)*, 2018.

[22] M. Hirth, T. Hossfeld, and P. Tran-Gia, "Anatomy of a crowdsourcing platform - using the example of Microworkers.com," in *2011 Fifth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing*, 2011.

[23] P. Ipeirotis, "Demographics of Mechanical Turk," *A Computer Scientist in a Business School [Blog]*, 03 2010.

[24] ——, "Mechanical Turk: Now with 40.92% spam," *Behind the Enemy Lines*, Dec. 16 2010 (accessed Feb. 25, 2020), https://www.behind-the-enemy-lines.com/2010/12/mechanical-turk-now-with-4092-spam.html.

[25] L. C. Irani and M. S. Silberman, "Turkopticon: Interrupting worker invisibility in Amazon Mechanical Turk," in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 2013.

[26] N. Jindal and B. Liu, "Opinion spam and analysis," in *Proceedings of the 2008 International Conference on Web Search and Data Mining (WSDM '08)*, 2008.

[27] JobBoy, "Jobboy," https://www.jobboy.com, accessed: Feb. 20, 2020.

[28] A. Kharraz, W. Robertson, and E. Kirda, "Surveylance: Automatically Detecting Online Survey Scams," in *2018 IEEE Symposium on Security and Privacy (SP)*, 2018, pp. 70–86.

[29] K. Lee, P. Tamilarasan, and J. Caverlee, "Crowdturfers, campaigns, and social media: Tracking and revealing crowdsourced manipulation of social media," in *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media (ICWSM-13)*, 2013.

[30] K. Lee, S. Webb, and H. Ge, "Characterizing and automatically detecting crowdturfing in Fiverr and Twitter," *Social Network Analysis and Mining*, vol. 5, 12 2015.

[31] S. Li, J. Caverlee, W. Niu, and P. Kaghazgaran, "Crowdsourced app review manipulation," in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '17)*, 2017.

[32] Y. Liu, Y. Liu, M. Zhang, and S. Ma, "Pay me and I'll follow you: Detection of crowdturfing following activities in microblog environment," in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI '16)*, 2016.

[33] A. Lua, "What counts as a video view on Facebook, Instagram, Twitter, and Snapchat? the buffer guide to video metrics," *Buffer Library*, 09 2017 (accessed Oct. 5, 2020), https://buffer.com/library/social-video-metrics/.

[34] Microworkers, "Microworkers," https://www.microworkers.com, accessed: Feb. 20, 2020.

[35] M. Motoyama, D. McCoy, K. Levchenko, S. Savage, and G. M. Voelker, "Dirty jobs: The role of freelance labor in web service abuse," in *Proceedings of the 20th USENIX Security Symposium*, 2011.

[36] A. Mukherjee, A. Kumar, B. Liu, J. Wang, M. Hsu, M. Castellanos, and R. Ghosh, "Spotting opinion spammers using behavioral footprints," in *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '13)*, 2013.

[37] B. Muller, "Link building & establishing authority," *Moz*, Accessed: Oct. 5, 2020, https://moz.com/beginners-guide-to-seo/growing-popularity-and-links.

[38] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock, "Finding Deceptive Opinion Spam by Any Stretch of the Imagination," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Jun. 2011, pp. 309–319.

[39] J. Ross, L. Irani, M. S. Silberman, A. Zaldivar, and B. Tomlinson, "Who are the crowdworkers? shifting demographics in Mechanical Turk," in *ACM CHI '10 Extended Abstracts on Human Factors in Computing Systems (CHI EA 10)*, 2010.

[40] S. Sannon and D. Cosley, "'It was a shady HIT': Navigating work-related privacy concerns on MTurk," in *Extended Abstracts of the 2018 ACM CHI Conference on Human Factors in Computing Systems (CHI EA '18)*, 2018.

[41] ——, "Privacy, power, and invisible labor on Amazon Mechanical Turk," in *Proceedings of the 2019 ACM CHI Conference on Human Factors in Computing Systems (CHI '19)*, 2019.

[42] P. Satya, K. Lee, D. Lee, T. Tran, and J. Zhang, "Uncovering fake likers in online social networks," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management (CIKM '16)*, 2016.

[43] J. Sprouse, "A validation of Amazon Mechanical Turk for the collection of acceptability judgments in linguistic theory," *Behavior Research Methods*, vol. 43, pp. 155–67, 11 2010.

[44] G. Stringhini, G. Wang, M. Egele, C. Kruegel, G. Vigna, H. Zheng, and B. Y. Zhao, "Follow the green: Growth and dynamics in Twitter follower markets," in *Proceedings of the 2013 Conference on Internet Measurement Conference (IMC '13)*, 2013.

[45] N. Su, Y. Liu, Z. Li, Y. Liu, M. Zhang, and S. Ma, "Detecting crowdturfing 'add to favorites' activities in online shopping," in *Proceedings of the 2018 World Wide Web Conference (WWW '18)*, 2018.

[46] Turkopticon, "Turkopticon," https://turkopticon.ucsd.edu/, accessed: Oct 6, 2020.

[47] G. Wang, T. Konolige, C. Wilson, X. Wang, H. Zheng, and B. Zhao, "You are how you click: Clickstream analysis for sybil detection," in *Proceedings of the 22nd USENIX Security Symposium*, 2013.

[48] G. Wang, C. Wilson, X. Zhao, Y. Zhu, M. Mohanlal, H. Zheng, and B. Y. Zhao, "Serf and turf: Crowdturfing for fun and profit," in *Proceedings of the 21st International Conference on World Wide Web (WWW '12)*, 2012.

[49] H. Xia, Y. Wang, Y. Huang, and A. Shah, "'Our privacy needs to be protected at all costs': Crowd workers' privacy experiences on Amazon Mechanical Turk," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. CSCW, Dec. 2017.

[50] L. Zhang and Y. Guan, "Detecting click fraud in pay-per-click streams of online advertising networks," in *2008 The 28th International Conference on Distributed Computing Systems*, 2008, pp. 77–84.

APPENDIX

A. Category Keywords

Table III provides the keywords that we used for category assignments. Words from a keyword phrase can appear in any position and order in the text. We caution against applying keywords directly as a classifier. As we discuss in Section III-C, our goal was not to produce a classifier, and we manually refined category assignments.

B. Active Campaigns over Time

To examine the long-term trend in active campaigns across platforms, we checked the total number of available campaigns on each platform again on August 10, 2020 and September 9, 2020 (in addition to February 19 and March 19). See Table IV.

C. Category Distribution by Platform

Table V shows the top five categories for each platform based on the mean per-day percentage of campaigns in each category. We exclude subcategories here. On each platform, the top five categories cover at least 80% of observed campaigns. On JobBoy, the top two categories were campaigns requesting personal information and campaigns that instruct workers to sign up for accounts, offers, and more. Minijobz shares these top two in reverse order. Like Wang [48], we suspect that sign-up campaigns were to boost the apparent user base for smaller services. SEO and social media were the most popular categories on Microworkers.

TABLE III
CATEGORY KEYWORDS FOR LEMMATIZED TEXT.

Category	Keywords
SEO (All)	*See subcategories*
—*Search term*	search [term, result, phrase, video, type, click]; click [result, organic]; quick search
—*Timed activity*	click [stay, minute, day]; [stay, visit, watch] [second, minute, day]; keep [second, minute, day, week]; [leave, delete] [day, week]
—*Backlink*	forum [link, website, post]; blog comment link; backlink site
Social media (All)	*See subcategories*
—*Engagement*	[facebook, instagram, linkedin] [like, share, comment]; twitter [like, share, tweet, retweet]; reddit upvote; snapchat like; youtube [like, thumbs, comment]; git star; medium [clap, comment, share]
—*Connections*	[facebook, instagram] [post, follow, friend]; linkedin [post, follow]; [twitter, snapchat, git, medium] follow; youtube subscribe; snapchat [follow, friend]
Personal information	[submit, confirm, enter, valid, register] [information, email, zip]; [submit, confirm, valid, register] number; [confirm, valid] name
Review activity	[positive, good, helpful, honest, relate, respectful, star] [leave, give]; [positive, helpful, honest, relate, respectful, relevant] write; positive answer; [positive, good, helpful, honest, relate, respectful, relevant] [comment, feedback]; [positive, helpful, honest, relate, respectful, star] rate; [positive, good, helpful, relate, leave, respectful, natural] review; natural comment; [positive, good, honest] experience; good rate; upvote [answer, question, review]
Sign up	sign up [account, complete, link, website, member, simple, registration]; register [account, complete, link, member]
Account use	account [need, require, must, have, log in]
Redirect	*Campaigns with URLs that lead to different landing pages or redirect through domains:* `appave.mobi`, `clkitgo.com`, `cpagrip.com`, `golead.pl`, `lnkclik.com`, `maxbounty.com`, `rotatemyurls.com`, `smrturl.co`, `viral481.com`, `xor-link.com`
Download/install	[install, download, add] [app, software, extension, android, apple, podcast]
Shopping	add [cart, product]; [wish, watch] list; [buy, purchase] [ringtone, product]; checkout; amzn continue
Spam filter	not spam; email short comment
Gift card	gift card; giftcard; gift-card; `onlinepromotionsusa com`; `promotionsonlineusa com`; `retailproductzone com`; `electronicproductzone com`
Account purchase	[email, mail] recovery; [create, open] [gmail, yahoo, outlook, hotmail, paypal] account
Ad activity	click [ad, banner]

TABLE IV
OBSERVED ACTIVE CAMPAIGNS.

Platform	Feb. 19	Mar. 19	Aug. 10	Sep. 9
JobBoy	201	161	168	127
Microworkers	1,193	582	945	768
Minijobz	25	19	18	12
Total	1,419	762	1,131	907

TABLE V
TOP CATEGORIES (EXCLUDES SUBCATEGORIES).

Platform	Category	Mean
JobBoy	Personal information	45%
	Sign up	36%
	Redirect	34%
	Download/install	14%
	Gift card	14%
Microworkers	SEO (All)	69%
	Social media (All)	57%
	Personal information	12%
	Sign up	7%
	Review activity	7%
Minijobz	Sign up	32%
	Personal information	32%
	Redirect	24%
	Download/install	22%
	SEO (All)	17%

D. Redirect Services

Table VI lists the ten observed redirect services behind redirect campaigns.

TABLE VI
OBSERVED REDIRECT SERVICES.

Redirect service	Campaigns
`cpagrip.com`	37
`viral481.com`	33
`maxbounty.com`	23
`lnkclik.com`	16
`clkitgo.com`	11
`smrturl.co`	10
`golead.pl`	2
`rotatemyurls.com`	2
`appave.mobi`	1
`xorlink.com`	1

E. Shopping Distribution

Table VII provides details for campaigns in the shopping category. For each combination of retailer, type of items involved (using retailer-provided categories where feasible), and requested action, the table provides the number of campaigns and the number of unique items specified by those campaigns.

TABLE VII
SHOPPING DISTRIBUTION.

Platform	Item Type	Action	Campaigns	Unique Items
Amazon	Home & Kitchen	Add to cart	11	36
	Health & Household	Add to cart	6	21
	Groceries & Gourmet Food	Add to cart	6	6
	Pet Supplies	Add to cart	5	16
	Tools & Home Improvement	Add to cart	5	13
	Sports & Outdoors	Add to cart	5	7
	Health & Personal Care	Add to cart	5	5
	Automotive	Add to cart	4	4
	Toys & Games	Add to cart	3	4
	Patio, Lawn & Garden	Add to cart	2	11
	Luggage & Travel Gear	Add to cart	2	2
	Office Products	Add to cart	1	7
	Multiple categories	Add to wish list	1	5
	Arts, Crafts & Sewing	Add to cart	1	1
	Clothing, Shoes & Jewelry	Add to cart	1	1
	Electronics	Add to cart	1	1
Steam	Games	Add to wish list	4	4
eBay	Home, Furniture & DIY	Add to watch list	3	3
iTunes	Entertainment	Purchase	2	2
Restaurant	Food delivery	Purchase	1	1