Beyond Monetary Incentives: Experiments in Paid Microtask Contests

OLUWASEYI FEYISETAN and ELENA SIMPERL, University of Southampton

6

In this article, we aim to gain a better understanding into how paid microtask crowdsourcing could leverage its appeal and scaling power by using contests to boost crowd performance and engagement. We introduce our microtask-based annotation platform *Wordsmith*, which features incentives such as points, leaderboards, and badges on top of financial remuneration. Our analysis focuses on a particular type of incentive, contests, as a means to apply crowdsourcing in near-real-time scenarios, in which requesters need labels quickly. We model crowdsourcing contests as a continuous-time Markov chain with the objective to maximise the output of the crowd workers, while varying a parameter that determines whether a worker is eligible for a reward based on their present rank on the leaderboard. We conduct empirical experiments in which crowd workers recruited from CrowdFlower carry out annotation microtasks on Wordsmith—in our case, to identify named entities in a stream of Twitter posts. In the experimental conditions, we test different reward spreads and record the total number of annotations received. We compare the results against a control condition in which the same annotation task was completed on CrowdFlower without a time or contest constraint. The experiments show that rewarding only the best contributors in a live contest could be a viable model to deliver results faster, though quality might suffer for particular types of annotation tasks. Increasing the reward spread leads to more work being completed, especially by the top contestants. Overall, the experiments shed light on possible design improvements of paid microtasks platforms to boost task performance and speed and make the overall experience more fair and interesting for crowd workers.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing theory, concepts and paradigms**; **Empirical studies in collaborative and social computing**; Social engineering (social sciences);

Additional Key Words and Phrases: Crowdsourcing, crowd computing, incentives, gamification, paid microtasks, contests, Markov chains

ACM Reference format:

Oluwaseyi Feyisetan and Elena Simperl. 2019. Beyond Monetary Incentives: Experiments in Paid Microtask Contests. *ACM Trans. Soc. Comput.* 2, 2, Article 6 (June 2019), 31 pages.

<https://doi.org/10.1145/3321700>

O. Feyisetan is now with Amazon Research, Seattle, USA (*work completed at the University of Southampton, UK*).

The work has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 732194 (QROWD), and from the UK EPSRC under Grant No. EP/J017728/2 (SOCIAM: The Theory and Practice of Social Machines).

Authors' addresses: O. Feyisetan and E. Simperl, Electronics and Computer Science, University of Southampton, UK; emails: seyi.feyisetan@gmail.com, e.simperl@soton.ac.uk.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2019 Copyright held by the owner/author(s).

2469-7818/2019/06-ART6

<https://doi.org/10.1145/3321700>

1 INTRODUCTION

More than a decade after the term was introduced by Jeff Howe in 2006 [33], crowdsourcing seems to be considered by many as a, if not “the,” universal means to solve virtually any kind of problem, online and offline, that requires sustained human involvement. We see it used to motivate employees to engage with less rewarding work routines, attract the best possible ideas to boost innovation, enhance artificial intelligence algorithms, and support ambitious social and entrepreneurial initiatives. Achieving a goal by collecting contributions from many individuals has a long tradition, far beyond the most recent developments in the digital world. In fact, some of the most successful exemplars of the “wisdom of the crowds” in modern times, for instance Wikipedia, predate Howe’s article. And yet, with the rise of social networks, smart mobile devices and online platforms, the phenomenon has found a new dimension in terms of scale and achievements—it routinely manages to mobilise very large groups of people in a relatively short period of time, helping organisations from tech and government to the military and marketing improve the ways they operate, decide, and engage with the world.

The crowdsourcing landscape is as diverse as its applications. It includes paid microtask platforms such as Amazon’s Mechanical Turk¹ and CrowdFlower (now Figure Eight),² alongside online labour marketplaces such as TaskRabbit³ and UpWork,⁴ and open innovation contests in the style of Kaggle⁵ and Innocentive.⁶ It features volunteer citizen science systems such as Galaxy Zoo⁷ and SciStarter,⁸ and games with a purpose (GWAPs) such as PhraseDetective⁹ and EyeWire.¹⁰ Putting aside the principled differences between these forms of crowdsourcing, which make them from the offset amenable only to specific types of problems, the success of any crowdsourcing endeavour will depend on the ability of the “requester”—the person or institution reaching out to the crowd—to benefit from crowd outputs and attract a critical mass of contributors. The frame conditions can be challenging as well: To be effective, crowdsourcing needs to provide a real alternative to existing solutions of the problem the requester is trying to solve, in terms of quality, timeliness, and costs. Getting this mixture right is not trivial, as it requires insight into a wide array of subjects from artificial intelligence and data analysis to user experience design and behavioural sciences [21, 22].

Aligning the motivations of the crowd with the goals of the requester and finding the right mix of rewards to drive purposeful crowd engagement are particularly critical. Existing crowdsourcing platforms already provide some level of support—whether paid microtasks, challenges, or citizen science projects, each of these forms of crowdsourcing makes some basic assumptions about what drives people to contribute to their “open call” [33], and offers sometimes bespoke incentives to encourage a certain style of behaviour. For example, people register on Mechanical Turk primarily for financial reasons [49, 50], then for learning, and possibly for fun. However, this is sometimes predicated on the task paying well. However, in contests such as the Netflix prize¹¹ or on TopCoder,¹² the crowd is mostly technology-savvy, and takes part mostly to gain a reputation, and

¹<https://www.mturk.com/>.

²<https://www.figure-eight.com/>. In the article, we will continue to refer to CrowdFlower, as the work was carried on before the rebranding as Figure Eight.

³<https://www.taskrabbit.com/>.

⁴<https://www.upwork.com/>.

⁵<https://www.kaggle.com/>.

⁶<https://www.innocentive.com/>.

⁷<https://www.galaxyzoo.org/>.

⁸<https://scistarter.com/>.

⁹<http://anawiki.essex.ac.uk/phrasedetectives/>.

¹⁰<http://eyewire.org/explore>.

¹¹<http://netflixprize.com/>.

¹²<https://www.topcoder.com/>.

then for the prize, or out of a desire to improve one's skills [1]. In citizen science, which is often devoid of financial rewards, the main motivator is a blend of learning, fun, and the wish to make a contribution to the world's knowledge [57, 67]. Designing a successful crowdsourcing campaign may therefore require different incentives, which respond to the different reasons that drive people to engage. This is not easy to achieve, as layering incentives on top of crowd motivations to change behaviour has proven to engender unexpected effects—for example, studies have shown that paying for work completed out of intrinsic reasons may make people contribute less, a phenomenon referred to as “crowding out” or “over justification” [28]. Similarly, incentivising people through gamification commonly fails to achieve significant, long-term results when there is a mismatch between the rewards induced through game elements and the needs, expectations and values of the crowd [62, 74].

In our work, we study the interplay between motivations and incentives in the context of *paid microtask crowdsourcing*. As financial payments are just one, albeit an important, factor in using crowdsourcing effectively, we explore the use of points, badges, leaderboards and contents to boost task performance and crowd engagement. To do so, we use our own experimental crowdsourced-annotation platform *Wordsmith*, which features paid microtasks at its core and allows requesters to layer different incentives on top of the financial remuneration and analyse their impact.

An initial study of *Wordsmith* and the effects of gamified task design on paid crowdsourcing was presented in Reference [25]. It showed that designing a playful interface for annotation tasks—as opposed to the functional style common to most microtask platforms—encourages workers to engage with the task more, independently of the actual monetary reward. In the study, we were consistently able to achieve significantly lower costs (5,708 unique annotations collected via the game vs. 111 unique annotations contributed in the control condition on CrowdFlower) and improve accuracy by almost 10% compared to the baseline. Through furtherance incentives, which were offered to workers when we predicted they were about to leave, we could convince workers to continue annotating, generating up to five times more unique annotations at a comparable level of accuracy. These effects were stronger when the incentives were personalised, with 77% of crowd workers deciding to complete one more image when confronted with a targeted furtherance incentives against 27% in the experiments using a randomly selected one. A follow-up study looked at the effects of social incentives [21, 22]—we asked workers to complete annotation tasks in pairs and to encourage their co-worker to continue so that both parties receive their payments. Again, the results were very encouraging. Adding a collaborative component to the default microtask model led to a 233% increase in task uptake, a 140% increase in task output and a 6% increase in inter-annotator agreement. We also demonstrated that social furtherance incentives create a win-win scenario for the requester and for the workers, by helping more workers get paid by re-engaging them before they drop out.

In this article, we focus on improving another critical success factor of crowdsourcing, delivery time. We developed an extension of *Wordsmith* that allows crowd workers to pick up tasks in near-real-time and compete. Workers were recruited from CrowdFlower, but instead of carrying out their annotation work at their own pace, they were asked to pay attention to the time-sensitive element of the tasks, and aim for a faster turnaround. Central to the competition was the leaderboard, which decides which workers will receive a payment for their work: unlike in most paid microtasks projects, which pay workers per unit of work accepted by the requester, in our study, workers competed for top positions on the leaderboard to get paid.

Crowd workers used *Wordsmith* to complete work for a fixed period of *time*. They were rewarded only if they achieved a high enough position on the leaderboard—in other words, workers competed against each other and only a share of them (captured by a so-called *reward spread*) received a payoff at the end of the contest. Leaderboard rankings were computed and updated in

real-time as a function of the number of tasks completed and a heuristic approximation of the quality of the participant submissions. The quality computation was based on previous work of ours, in which we studied different approaches to crowdsource the same annotation task, using standard paid microtask platforms [20, 23].

The main intuition behind the contest model was that by designing the crowdsourcing exercise as a live contest, which is supposed to be completed in a relatively short period of time, we would create an environment in which results would be delivered quickly and with accuracy. In addition, as we did not pay all workers upfront or merely for being available, we could keep the overall costs lower than in other real-time crowdsourcing approaches, which often reward mere availability as opposed to work being done. In this context, we hypothesised that the number of workers rewarded impacted the output volume.

To test this, we ran experiments with three sets of reward spreads (top worker; top 5 workers; and top 10 workers). The experiments confirmed our initial assumption: the contest model resulted in more annotations at a rate that was twice as fast as the control condition reported in our previous work [20]. Furthermore, increasing the reward spread led to an increase in task output. Rewarding more workers reduced the rate of worker attrition and kept more workers engaged.

These results could help in finding the balance between theoretical guarantees with empirical evidence to select appropriate reward spreads, while scaling to data streams comparable to real-world situations.

Summary of contributions. Our platform Wordsmith offers a lens through which we can better observe the potentials of different incentive mechanisms in paid microtask crowdsourcing. Our main observations are the following: While *financial payments* remain important, *contests* can be extremely effective in encouraging workers to volunteer to carry out work at higher speed. *Leaderboards* make workers more aware of their performance—they drop off when faced with a potential loss of utility, i.e., when they believe they might not be among the winners of the competitions and hence miss on the payments, but the faster-paced design does not have significant effects on output quality or user experience.

The experiments shed light on possible design improvements of paid microtasks platforms to boost task performance, and make the overall experience more fair and rewarding for the workers. While we are not necessarily arguing for fully fledged contest-based microtask platforms, considering specific gamification elements that are widely discussed in the literature [40, 74] is worth further investigation. This is important not just for purely utilitarian motives on the side of the requesters but also in the context of the ongoing debate on ethical and fair crowdsourcing [34]. In the same time, competitions, which are widely used in other forms of crowdsourcing (for example, in citizen science [67] and in challenges [66]) offer interesting alternatives in situations when having a faster turnaround is important.

Previous work has approached such questions mostly through studies of crowd motivations [35], discussing the rich repertoire of extrinsic and intrinsic reasons that drive people to complete microtasks. Our experiments quantify some of their findings [20, 22]. We deliberately chose a task that is well-known in the crowdsourcing literature [30], as we were aiming for microtask designs that were only minimally influenced by interface or quality control aspects. For the same reasons, we opted for average market prices to reward participation; lower pays would have been less attractive (and unfair) for workers, higher ones might have appealed to people who were exclusively drawn to financial incentives. We believe more research needs to be done to build microtask platforms that are crowd-, as opposed to requester-centric, and that reflect and support the values and motivations of the crowd as an integral part of their functionality. Our research gives evidence that such efforts could be beneficial for both workers and requesters—points, leaderboards and

contests can be linked to intrinsic motivation. Workers not only completed more tasks than required, but some of them reported positively about the experiments on a community forum—for example, one post read “*Hello everyone! Lately I’m hooked on the multiplayer tasks, waiting for 100 people to connect,*” while another one claimed: “*Hit the top 10 today. I will hunt this problem again.*”

Crowdsourcing, though riddled with its challenges from the technical to the ethical, is here to stay. Even as it morphs continually to the tune of academic and economic forces, it remains a force for good with groundbreaking discoveries being made daily directly and indirectly through the wisdom of the crowd [32]. Understanding the impact of incentives brings us one step closer to designing better, more responsible crowdsourcing systems. This is vital not just for research purposes but also because of the influence crowdsourcing has on fields that have come to rely on it, including artificial intelligence, and those that draw inspiration from it, such as the so-called “gig” economy.

Previous publications. This article builds upon previous work of ours, in particular on References [20, 24, 25]. References [20, 25] used earlier versions of Wordsmith to carry out their experiments. In Reference [25], we presented a first version of Wordsmith and initial experiments into the effects of gamified task design on paid crowdsourcing. In References [20, 24], we aimed to understand how to collect Twitter entity annotations via paid microtasks effectively. All works implemented a task model where workers complete tasks individually and are rewarded if they provide useful answers. In this work, we used the same Wordsmith platform and Twitter entity annotation task; however, we introduced contests as a new incentive layer on top of the base financial payouts.

2 BACKGROUND

In this section, we briefly review some of the most relevant prior work pertaining to maximising the effectiveness of incentivised crowdsourcing. In particular, we focus on real-time crowdsourcing approaches that aim to generate results faster; and more generally, on methods that optimise some aspect of crowd performance, be that by offering bespoke incentives, by assigning tasks to those workers who are more likely to do them well and so on. As much of this background literature is inspired by and explained using theories of human motivation and incentives design, we start with a synopsis of fundamental work in this space.

2.1 Incentives in Paid Microtask Crowdsourcing

2.1.1 Motivation Theory. Theoretical frameworks such as the Self Determination Theory (SDT) [15] make the distinction between intrinsic and extrinsic motivation to study the reasons why individuals decide to contribute to a task. Intrinsic motivation covers tasks that are perceived as rewarding in themselves, and is seen as responsible for engagement in activities such as socialising and participating in volunteering projects [67]. However, extrinsic motives are related to factors that are not inherently related to the actual task but are appealing for some external reason such as status or financial payments.

The SDT theory has been applied to motivation studies on paid microtask platforms. Reference [35] mirrors the motivation types as *fun* (intrinsic) and *money* (extrinsic) in a survey of incentives on Mechanical Turk. Similarly, Reference [48] describes incentives as either appealing to *love*, *glory*, or *money*. More relevant to Wordsmith, SDT has also been applied extensively in gamification research [17, 63]; collaboration and online communities [56, 57, 67]; and contests [3].

2.1.2 Monetary Incentives. Money is a natural incentive for carrying out work, hence paid microtasks remain one of the most prominent forms of crowdsourcing [27]. Research has shown that increased payment leads to faster task completion, but not to higher quality [50]. Several other

factors play a role in obtaining useful task results: bonuses, worker perception, and the variation of payment sizes across tasks [18, 50]. The drive to engineer optimal quality, speed and volume of work with minimal financial payments has made paid microtasks the subject of numerous discourses on the ethics of compensation [34]. While monetary payments can be seen as an extrinsic motivator, other intrinsic factors have been known to contribute to sustained participation in paid crowdsourcing settings, including intellectual curiosity [46]; the inherent nature of the task [27]; and interaction with other workers [36]. One of the questions we are trying to answer with Wordsmith is: *How do we expand the currency of transactions on paid microtask platforms to transcend monetary payments and encompass tasks that have an intrinsic appeal?* We presuppose that framing tasks this way would be more attractive to workers, cost less for requesters, and make workers perform better. This led us to study other incentive mechanisms in our earlier works [20, 22], and the idea to leverage contests as a new incentive strategy.

2.1.3 Gamified Incentives. Gamification is the use of game design in non-game contexts, with the aim to achieve the effects of fun and engagement that are created by playing a game [74]. It often involves adding game-like rewards to non-game activities, including social elements to encourage team play or contests. We understand gamification in a broad context, comprising systems that build a complete game narrative around a task (e.g., FoldIt and EyeWire) [32]; employ tactics such as micro-diversions to fend off boredom [13]; further a noble cause or stoke curiosity [46]; or engineer game mechanics (such as points, badges, levels and leaderboards) into tasks [17, 39]. In a work context, each of these classes of systems has been shown to improve productivity, drive engagement and effectiveness, and therefore reduce costs.

Gamification practices have raised concerns about their potential to undermine innate intrinsic motivation via the crowding-out and over-justification effects discussed earlier—although this has been shown to be task-dependent [51, 63]. We also know that gamification can be more effective in collaborative settings; in other work of ours [22], we designed a multiplayer Wordsmith extension that required consensus among pairs of paid crowd workers. This led to better outcomes when compared to the baseline, which was built around completing the same tasks individually [22]. Similar findings were reported in the context of GWAPs [70], which often rely, unlike our work, on voluntary participation.

2.1.4 Contest-based Incentives. Paid microtask crowdsourcing has traditionally been approached as an individualistic endeavour. Workers complete their tasks without depending on, or interacting with others. Other forms of crowdsourcing, such as volunteer citizen science, have experimented with models that encourage the crowd to discuss, exchange ideas, collaborate on tasks, or compete to improve their standing in the community [58, 67]. In this context of paid microtasks, we seek to apply contest-based incentives to improve the delivery time of crowd responses.

To carry out crowdsourcing in a contest setting, timely worker recruitment and task uptake are critical. This comes as a result of scenarios with near-real-time constraints—for example, when crowdsourcing is applied to fact-checking in news media, disaster management, or live data feeds. However, most solutions proposed in the literature are often too costly for large numbers of microtasks (see also Section 2.3 below for an overview of near-real-time crowdsourcing research). Contests recruit crowd workers who are attracted to the thrill of “glory” and competition to carry out tasks quickly. This model has been seen on software coding platforms such as TopCoder [1], where over a million workers vie to complete challenges, or in citizen science [58]. Contests leverage on the urgency that comes from a fixed completion time-frame, and the satisfaction brought by a sense of winning. They are primarily deployed when the requester seeks one best or final answer (as opposed to an aggregation of worker results, which is more typical in microtask crowdsourcing), or to increase participation. For example, the Netflix \$1 million challenge to build a better

recommendation algorithm [3] falls in this “single best result” category. EyeWire, a game with a purpose that creates a visual map of the human brain, routinely uses different kinds of challenges to increase the volume of tasks completed and keep the crowd engaged [58]. Remuneration in contests range from a winner-takes-it-all model, which compensates only the best participant, to more relaxed models that pay contributors who make submissions above a certain threshold, with several theoretical frameworks set up to determine the optimal allocation of prizes [53]. In this article, we designed our contest to cover tasks that required an aggregation of results (as opposed to a single best response), while obtaining responses as quickly as possible. This type of approach is relevant in domains such as disaster relief and real-time visual aids for blind people [7].

The literature on contests goes well beyond the focus of our work, which is paid microtasks. The history of contests dates back to as far as 1714, when the British Parliament ran one to determine the longitude at sea to within half a degree [53]. Several theoretical studies have looked at optimal contest design [11] and the optimal allocation of prizes [53]. A survey of experimental research of contests is available from Reference [14]. Researchers have investigated how contests unfold on existing platforms such as TaskCN [47], TopCoder [1], or EyeWire [58]. Empirical studies have shown the effects of increased payoff as an indicator of contestant performance. Further on, Reference [8] considered over 9,000 contests hosted on the same platform to understand the effect of participant numbers on the performance of individuals. A second group of empirical work focused on bespoke experimental setups for crowdsourcing. For example, Reference [60] looked at the effect of varying monetary schemes and information policies in individual contests, while Reference [61] explored the same problem alongside team formation strategies. Both papers bear similarities to our scenario, in which contestants strategically decide whether they continue to take on more tasks or leave Wordsmith. However, they do not propose ways to predict worker exit to optimise task completion and delivery. Reference [55] introduced a duration model, which we also attempt to analyse in our work, which shows the length of candidacies and factors associated with candidate exits.

Finally, our work builds on literature that studies the war of attrition. In contests, each participant enters knowing their own skill and costs, but not those of the other contenders. Participants consider dropping out when they learn their opponents’ strengths and discover that staying would be unprofitable. A theory of how this phenomenon operates in duopolies was presented by Reference [29]. Reference [41] presented its relation to an all-pay auction while, while Reference [55] discussed attrition and exit in political primaries. Reference [54] looked at contests with exits where contestants have the option to drop out or not to participate with the introduction of costless punishments. Since our objective was to maximise the total utility generated in real-time, we did not use punishments, which some early pilots we carried out revealed to increase the attrition rate.

2.2 Making Microtask Crowdsourcing More Effective

Several descriptive frameworks have been proposed in the literature to capture the nuances of incentives engineering beyond simplistic “*fun or money*” considerations. Some of these include *MICE* (Money, Ideology, Coercion, Excitement) [9]; *RASCLS* (Reciprocation, Authority, Scarcity, Commitment, Liking, Social Proof) [9]; and *SAPS* (Status, Access, Power and Stuff) [74]. The latter is intended to represent a system of incentives from the most desired to the least desired, and from the cheapest to the most expensive. We adopted this framework in Wordsmith.

Mechanisms for the effective allocation of incentives have been studied in market and auction platforms, wireless and peer-to-peer networks and corporate organisations. In the context of crowdsourcing, several studies have been carried out that apply *game-theory* techniques to incentive design [71, 73]. These papers focus on financial incentives and a premise of inter-player

strategy dependency. Not all crowdsourcing tasks can be modelled in this way; we adopt a probabilistic approach based on prior player behaviours to predict appropriate incentives beyond the purely financial (see also Reference [20]). Similar techniques are used for various purposes in crowdsourcing design, in particular to inform the assignment of tasks to workers or to predict task completion [16, 65].

A large body of work has been dedicated to task and workflow design, as well as quality control (see, for instance, Reference [52] for a good compilation). We take their findings into account when implementing the basic CrowdFlower interfaces, checking for spam and validating and aggregating the results.

2.3 Improving Delivery Time

Longitudinal studies of crowdsourcing marketplaces such as Mechanical Turk [18] reveal how microtasks have evolved to support scenarios in which work cannot take days or hours, but must be delivered in seconds and under [4]. For near-real-time behaviour, two components are critical: prompt recruitment of workers and fast completion of tasks. In this section, we compare our approach with previous research in microtask crowdsourcing, which has proposed improvements to these components.

2.3.1 Timely Worker Recruitment. Several models have been put forward to ensure timely availability of crowd workers. For example, Adrenaline [4] employed a retainer model, while in Viz-Wiz [7] workers signed up in advance and were kept engaged, at a cost, until the work arrived. Other approaches relied on very large crowds [42], using queuing theory or predictive recruitment to model workers' arrival and task assignment [5], or repeatedly posting the tasks to make them more visible to workers [7]. In our experiments, we incorporated a combination of these ideas to generate our on-demand crowd: We repeatedly posted tasks, recruited in advance, requested for a larger crowd than required, and used an audio alert to prompt workers when the contest started. At the same time, by virtue of our contest model, we rewarded only a small share of the workers, hence keeping the costs down.

2.3.2 Timely Task Completion. Some of the techniques found in the related literature include: *rapid refinement*, which dynamically narrows the search space following signs of worker agreement [4]; *stream parallelism*, which splits the task into sub-tasks and assigns them to different pools of workers in parallel [7]; *temporal division*, which breaks the task down into time slices as more work becomes available [45]; *warping time* [44], which allows workers to listen to audio streams at reduced speeds; and a *Map Reduce* alike to organise complex workflows [38]. Our model uses stream parallelism to assign work to several groups of workers, and variants of temporal division and time warping to display tasks to each overlapping batch of workers for an extended period of time. Combined with the mixed cardinal-ordinal contest, these techniques create a powerful mechanism to solve tasks twice as fast as the evaluation baseline (see Section 7.2).

3 CONTEST-BASED ANNOTATION TASK

In our experiments, we used a text annotation task model in relation to real-time crowdsourcing. For this reason, the task consists of a total of n posts, $P = \{p_1, \dots, p_n\}$, each containing m entities $E = \{e_1, \dots, e_m\}$ to be annotated, where $m < M_i! + M_i$ and M_i is equal to the number of text tokens in post p_i . The posts arrive at a constant rate λ and each has a processing rate of μ . There are n workers in a pool to serve the task queue such that, to keep up with the requests, the ingress load (task intensity) $L = \lambda/\mu$ stays less than the number of workers n , i.e., $L < n$.

Tasks that are not solved are dropped of the queue as opposed to being kept indefinitely in the buffer [5]. They are solved using a first-in-first-out scheduling and processing scheme, in which

already recruited workers are sought to carry out new tasks (as opposed to recruiting additional workers for new tasks). Therefore, the requester is looking for an optimal processing rate μ , and needs to keep workers motivated to carry out as many tasks as possible.

3.1 Operationalised Variables

Our task model had one independent variable: the reward spread, and one dependent variable: the task output.

Independent variable: *reward spread*

This represents the total number of prizes allocated by the requester to participants in the contest run. For example, a reward spread of 3 indicates that the top 3 annotators would get paid. The reward spread was instantiated as a leaderboard that ranked users as a function of the number of high quality annotations they make.

Dependent variable: *task output*

This represents what the requester seeks to optimise for by varying the reward spread. It is the total number of accepted (high quality) posts labelled by crowd workers within the annotation session. A proxy for measuring the quality of annotations was pre-computed using a parts-of-speech (POS) tagger to estimate the expected number of entities.

3.2 Task Elements

Requester. The requester asks the crowd to complete a series of tasks in real-time. The requester needs to determine the experiment setup in terms of: (i) the number of contestants; (ii) the number and size of prizes; and (iii) the contest constraints to maximise the effort exerted by all contestants $\sum_{i=1}^w \epsilon_i$. This is different from contests such as the Netflix Challenge [3], where the principal's objective was to elicit a single best response to a task.

The requester does not only desire to maximise the total exerted effort but also to maximise some utility function of output quality $\sum_{i=1}^w q_i$. The requester therefore needs to maintain incentives for highly skilled workers, and motivate low skilled workers to exert more effort while adjusting the prize spread.

The requester defines: (i) a completion time constraint T , which depends on the number of posts n and their arrival rate λ ; and (ii) a quality constraint Q , which denotes the minimum number of labels expected from each worker to be eligible for payment. The latter is essential in hybrid tasks; for example, the task might have been pre-labelled by a machine to determine the probable number of named entities (this serves as the quality constraint Q), while the crowd workers identify those entities and type them [20].

The requester is able to observe the baseline quality of each worker's output, based on the pre-computed number of entities and the number of entities submitted by the worker. The requester is then able to use this information to construct the contest by assigning a quality score to every crowd answer—therefore, contestants are ranked not only based on their effort (number of posts annotated) but also on the quality of their output. The reward mechanism awards a prize A_j to a worker w_j within a reward spread (e.g., the worker was within the top 5 or top 10).

Workers. There is a set of n workers, $W = \{w_1, \dots, w_n\}$ participating in the contest, each with the ability to carry out entity annotations. Each worker $w \in W$ has a private skill level ζ_i (also known as expertise or ability), and for each post in an annotation task, chooses to exert a level of effort $\epsilon_i \geq 0$. The skill level is drawn independently of other workers from the interval $\zeta_i \sim [0, 1]$, $\forall w \in W$, according to a distribution function F with density $f(\zeta) = dF(\zeta) > 0$.

The effort exerted is drawn from the interval $\epsilon_i \sim [0, \epsilon]$, in which the maximum effort expendable is constrained by the running time t of the contest, which in turn is a function of posts per

unit time and total number of posts. The quality q_i of each worker w_i is determined by: (i) the skill level ς_i ; (ii) the effort exerted ϵ_i ; and (iii) a requester variable δ_i . The requester variable δ_i is a function of the requester's review process and perception of quality, in comparison with the worker's internal tagging bias, which is markedly present in human judgement tasks. This value is constant across annotation posts $\forall p_i \in P$; therefore, two workers exerting the same effort to annotate the same post would differ only on their skill, since the requester's variable is constant for that post. The quality of a submission is thus given as: $q_i = \varsigma_i \epsilon_i + \delta_i$. In our experiments, the requester's variable was a measure of results in a pre-computed gold standard set [20].

Each worker w_i seeks to maximise their expected utility. This depends on the number and value of prizes, and on the number of contestants and value of their efforts. A worker's utility U_i is given by $U_i = V_j - c(\epsilon_i)$, if the worker w_i wins prize V_j , or by $U_i = -c(\epsilon_i)$ otherwise, where V_i is one of k prizes to be awarded by the requester, and $c(\epsilon_i)$ is the worker's cost function, which is a strictly increasing function dependent on exerted effort where $c(0) = 0$. Each prize V_j above point k is positive, or zero otherwise. We did not model negative rewards (punishments), as they did not seem to have the desired effects in early pilots we carried out.

During the contest, crowd workers are shown a list of tasks that arrive in near-real-time (in our experiments, we used Twitter posts, but they could equally be images as in Reference [7] or other objects). The workers are to annotate as many of them as possible within a time window. The incentives mechanism awards a prize A_j to a worker w_j within a reward spread (for example, the worker is within the top 5 or top 10).

Exit. A worker would always seek to maximise their expected utility given the number and value of prizes; the number of contestants; and the value of their efforts. In our experiments, it was possible for workers to view their ranked position in real-time with respect to their closest contenders using a k -neighbours leaderboard, as presented in the medium information policy contest strategy by Reference [60]. A worker far outside the reward spread might inadvertently decide to exit the contest to avoid further loss of utility. The worker close to the reward spread might, however, decide to remain in the contest in the hope of displacing a close contender.

4 CONTESTS AS CONTINUOUS-TIME MARKOV CHAINS

4.1 Model

We modelled the contest-based annotation task as a continuous-time Markov chain (CTMC). To understand CTMCs, we first define some key concepts.

A *stochastic process* is defined as a collection of random variables. It describes the evolution in time of a random phenomenon. Therefore, we say that the stochastic process is indexed by time, i.e., the random variables X for each time point t is given by $X = \{X(t) : t \in T\}$. When $T = \mathbb{N}$, it is a discrete-time process; however, when $T = [0, \infty)$, it is a continuous-time process.

The *Markov property* refers to the *memoryless* property of a stochastic process. A stochastic process $X(t)$ has the Markov property if the conditional probability distribution of its future evolution depends only on its current position, not on how it got there; i.e., given $X(t) : t \geq 0$ with state space S , we say it has the Markov property if

$$P(X(t) = j | X(s) = i, X(t_{n-1}) = i_{n-1}, \dots, X(t_1) = i_1) = P(X(t) = j | X(s) = i), \quad (1)$$

where $0 \leq t_1 \leq t_2 \leq \dots \leq t_{n-1} \leq s \leq t$ is a sequence of $n + 1$ times and $i_1, i_2, \dots, i_{n-1}, j \in S$ are states in the state space. Therefore, a continuous-time stochastic process $X = \{X(t) : t \geq 0\}$ is called a continuous-time Markov chain if it exhibits the memoryless Markov property.

Revisiting our contest-based annotation task, we assume that the n Twitter posts arrive according to a Poisson process with rate μ . This Poisson process is a continuous-time Markov chain with

state space $S = \{0, 1, 2, \dots\}$ representing the worker annotations. Therefore, at any time t , a worker i is in a state of annotating the j th tweet. The time it takes for the worker i to annotate this tweet is represented as an exponential random variable, which transitions the worker to state $j + 1$.

The question then is: What “explains” the rate at which a worker annotates each tweet? In our modelling in this section, and in subsequent empirical experiments, we evaluate one variable to explain the annotation rate. We have earlier described this as the *reward spread*, which governs, at any time t , if a worker i is *currently* eligible to receive a payment based on their leaderboard rank.

To model our contest-based annotation task as a continuous-time Markov chain, we define the following variables:

Observed variables:

- $S_i(t)$: the number of annotations of worker i in finite time t
- t_j^i : the time of the j th annotation of worker i
- τ_j^i : the annotation time for worker i after the j th annotation

Hidden variable(s):

- θ_i : the annotation rate parameter for worker i

The time τ_j^i that a worker i spends on annotating a tweet is represented as an exponential random variable based on our Markovian assumption that this depends on the preceding time:

$$\begin{aligned} \tau_j^i \sim P(\tau^i \mid \{S_k(t_k^i)\}_k; \theta_i) &= \exp(\tau^i; \lambda \cdot f(\{S_k(t_k^i)\}_k; \theta_i)) \\ &= f(\{S_k(t_k^i)\}_k; \theta_i) \exp(f(\{S_k(t_k^i)\}_k; \theta_i)) \tau_j^i. \end{aligned}$$

Furthermore, the time of the next annotation, and the next state can be expressed as

$$t_{j+1} = t_j + \tau_j^i; S(t_{j+1}) = S(t_j) + 1.$$

The parameter λ describes the annotator “behaviour” within the *reward spread*. This biases workers to either exert more or less effort based on their rank, i.e., if a worker is currently eligible to be paid, then they might either choose to exert more effort to extend their lead, or they might relax, since they are already in the top ranks. Similarly, a worker outside the reward spread might either give up or exert more effort to displace one of the top workers. We modelled λ using the gamma distribution $f(x; \alpha, \beta)$, and half-normal distributions $f_Y(y; \sigma)$, where $Y = |X|$ and $X \sim \mathcal{N}(\mu, \sigma^2)$. The half-normal distribution biases our belief that some workers will increase their effort, especially when they are outside but close to the reward spread. Alternatively, a normal distribution could be selected to model the possibility of workers either increasing or reducing their effort when they are eligible to be paid. We opted for the simpler (half-normal) distribution as designing a richer model of λ that fully captures the ability of humans to make rational choices, which maximises their utility will make direct analyses more complex. Nevertheless, we start off with a model that can be easily extended to more complex scenarios:

$$\lambda \sim \begin{cases} \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}, & \text{if worker is eligible to be paid,} \\ f(x; \alpha, \beta) + \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(-\frac{y^2}{2\sigma^2}\right), & \text{otherwise.} \end{cases}$$

Therefore, simplifying and expressing this for a single worker:

$$P(\{t_j^i\}_{j=0}^{N_i} \mid \{S_i(t_j^i)\}_{j=0}^{N_i} \mid \theta_i) = \prod_{j=0}^{N_i} f(\{S_k(t_j^i)\}_k; \theta_i) \exp(f(\{S_k(t_j^i)\}_k; \theta_i)) \tau_j^i.$$

Fig. 1. State transitions.

Computing this for all the workers, we get

$$P(T, S \mid \theta_i) = \prod_{i=1}^n \prod_{j=0}^{N_i} f(\{S_k(t_j^i)\}_k; \theta_i) \exp(f(\{S_k(t_j^i)\}_k; \theta_i)) \tau_j^i.$$

Expressing this as a log loss for training a learning model, we obtain

$$-\log P(T, S \mid \theta_i) = \sum_{i=1}^n \sum_{j=0}^{N_i} -\log f(\{S_k(t_j^i)\}_k; \theta_i) + f(\{S_k(t_j^i)\}_k; \theta_i); \tau_j^i. \quad (2)$$

Figure 1 highlights the state transitions for workers across the contest time steps with respect to their annotation rates. Figure 2 represents a graphical model that presents the variables τ , λ , f , and θ in our model. The arrows describe how each variable influences or affects another. Therefore, obtaining the values τ_j^i (i.e., for each worker i at every annotation point j) helps us predict how much time the annotator will spend in the entire contest.

5 THE WORDSMITH PLATFORM

Having laid the theoretical models underpinning our experiments, we now give a general overview of the Wordsmith platform and elaborate on those components that are relevant to our contest experiments.

5.1 Overview

Wordsmith is a platform for carrying out paid microtask crowdsourcing; Figure 3 depicts a screenshot. It supports a range of annotation tasks, including text and image annotation. Further on, Wordsmith features multiple task, feedback and interface elements that allow for single-worker, multi-worker and contest-based task completion.

The general task annotation works as follows: Workers are recruited from an existing microtask platform, for instance CrowdFlower, and redirected to our platform; the entry on CrowdFlower explains the task (e.g., identify people in tweets); the requirements (e.g., annotate at least ten images);

Fig. 2. Plate notation for CTMC contest-based annotation task.

Fig. 3. Wordsmith start interface. On the left-hand side, the worker sees their current level, the number of tasks completed, the number of points obtained, and their badges. On the right-hand side, they can consult the leaderboard and a feed of activities of other workers who engaged with the task.

and the financial remuneration. It also includes a tutorial consisting of screenshots of the task interface on Wordsmith to give workers a step-wise idea on how to approach the task effectively.

Requesters specify the conditions for workers to receive their payments—for instance, a minimum number of annotations being carried out. However, workers are generally allowed to carry out as many tasks as they wish, beyond the task threshold, which guarantees their payment. The requester can hence take advantage of the fun factor and game-like immersion that comes from completing the task in a gamified setting. Workers are also allowed to skip tasks—they can still receive their reward as long as overall they meet the conditions of the requester.

Requesters can also define game levels associated with an increasing volume of work to be completed or, if possible, with work of increasing difficulty. They also determine which worker activities are rewarded with gamification elements, including points, badges, and, as in Reference [20], additional payments.

5.2 Contest Design

The contest setting on Wordsmith allowed for multiple workers to connect to the system at once and carry out annotations simultaneously.

The current implementation of contests in Wordsmith builds upon the crowdsourced NER work, which was presented in References [20, 23]. In that work, we discussed the basic design of a text annotation microtask for Twitter posts, while here the focus is on the specific challenges arising from the fact that the annotations need to be delivered under significant time constraints. While the contest model was tested for NER on microposts, it can be used for other types of labelling tasks that arrive in a real-time feed, with or without predefined label categories, for example for images like in Reference [7]. For other text streams, the system would need first to divide the text into smaller chunks, which can still be treated by workers as microtasks in terms of the length of the text. Reference [6] offers interesting insights into how to process larger text documents via paid microtask crowdsourcing.

5.2.1 Input and Output. The system takes in a raw input of streaming posts and performs an initial sequence of processing on it. For our study, we focus on filtering out non-English tweets using the language tag of the incoming tweets. We then carry out part-of-speech tagging to recognise tweets with proper nouns. This is used to build a pseudo-quality score for each annotation (presented earlier as the requester's variable δ_i); i.e., if our POS tagger detects two different contiguous proper noun sets, then we can expect an annotation result of at least two entities (although this does not hold strictly if there are no proper nouns, e.g., entities might be recognised as noun phrases by some taggers).

The system outputs a processed stream of English tweets ($tweet1, \dots, tweetN$), where each tweet is represented as a tuple containing a reference identifier, the tweet string and the associated requester's variable δ_i . Each tweet in the stream advances in linear discrete time at a constant rate, with each time point represented as unique integer value in seconds (although, worker annotation and exit is represented in milliseconds).

5.2.2 Temporal Division and Stream Parallelism. The stream is bucketed into distinct time intervals using *windows*. A window consists of a constant number of microtasks emitted per unit time, which may or may not have been built over a buffer depending on varying throughput levels. Within each window task slice, microtasks are clustered and parallelized to different workers. Each cluster is a task unit consisting of a list of objects that is allocated to each worker for annotation. This follows a Map Reduce paradigm wherein each worker has a small task unit to solve, which is recursively built up to the final solution for the requester. The Map process involves local processing on individual nodes (annotating workers), while the Reduce process involves the merging of results to select best responses for overlapping task annotations.

5.2.3 Contest Interface. The contest interface consisted of a central annotation panel, in which the workers saw the current list of text snippets as described in References [20, 23]. After selecting the piece of text to analyse, the workers could either annotate the entities using predefined categories, mark the post as having no entities, skip the task all together, or go back to the list. Figure 4 shows a screenshot.

Fig. 4. Wordsmith contest interface. The central pane contains a microtask feed, which is updated according to the parameters fixed by the requester. Each task requires the player to identify instances of people, locations, and organisations. On the right-hand side, the players can see how they perform compared to others.

Each worker received a baseline score x for annotating a tweet with an arbitrary number of entities and a higher score $5x$ for correctly annotating the pre-computed number of entities (based on the gold standard tags). These figures were drawn from a series of observations and early pilots, which also ruled out the use of negative scores and qualifying questions, as both led to a very sharp rate of exits. As noted earlier, workers had access to a k -neighbours leaderboard, which lists contenders immediately above or below.

5.2.4 Requester Parameters. Wordsmith affords the requester a number of configurable parameters including:

- **Number of workers W :** count of those who could connect to the platform. All the workers had to be connected before the task starts.
- **Leaderboard:** requesters can choose between listing the *top- k* workers; or displaying the k contenders above and below the worker.
- **Number of tasks P :** This represents the total number of tasks that are to be processed, each consisting on n objects to be annotated.
- **Window size w :** This is the number of tasks that is sent out to a number of workers per time slice (see below).
- **Task unit time μ :** this is the delay time, for which a list of tasks remains available for annotation to a worker before the next set of tasks arrive.
- **Task arrival rate λ :** this variable measures the number of tasks, which are channelled to the platform per unit time.
- **Task intensity L (or ingress load):** this is computed as $L = \lambda/\mu$, based on the task arrival rate λ and the agent processing rate μ .
- **Task unit size:** represents the number of tasks that a worker actually sees on screen at any given time, which is a fixed percentage of the window size w .

- **Total task time T** : computes how long the contest would take as a function of the number of tweets and the window size: $T = P\mu/w$.
- **Reward spread s** : the reward spread denotes the number of workers to be paid in the current contest.

5.2.5 Warping Time. Warping time is a strategy, in which a workers task slice in a real-time assignment is deliberately slowed down to afford for maximal worker cognition in undertaking the required task. For example, Reference [44] used time warping to slow down audio playback so crowd workers could effectively transcribe a given portion of speech. This was recursively done for each worker, after which the individual results were successively merged to create a single result. Following our stream parallelism approach to dividing up the incoming microposts, a window of 200 tweets was presented to 20 workers.

We adopted the approach by Reference [44] denoting: (i) an in-period P_i where the annotation stream for a worker group comes in; (ii) a speed reduction rate rr ; and (iii) the compensating out-period P_o where the worker group WG_i rejoins the live real-time stream. We used a speed reduction rate rr of 10, i.e., the 20 workers experienced a streaming rate of $1/rr = 0.10$ (i.e., one-tenth speed, 10s rather than 1s) while annotating. During these 10s, the workers were presented with static list of tweets, from which they could select individual entities to annotate. The stream of tweets in the buffered out period was then emitted at a speed of

$$\frac{N - 1}{N - rr}. \quad (3)$$

In our experiments, we made certain assumptions, which would be handled differently on a live data feed. For example, all members of a worker group (20 workers in our case) were presented with the annotation tasks at the same input period P_o , which was a function of our streaming approach. We re-purposed existing datasets from literature into a streaming API. In actual practice, each worker would have a unique input period P_o similar to Reference [44].

5.2.6 Task Allocation. There are several approaches to distribute the incoming stream of tasks to the available workers in parallel. In a *round robin* assignment, each bin would be sequentially assigned from the window of current annotation tasks to the next available worker, e.g., in our experiments, we assigned each worker 10 tasks from 1 out of 20 bins derived from the streaming window of 200 tasks.

6 EXPERIMENTS

In this section, we describe how we carry out empirical contest-based crowdsourcing experiments using Wordsmith.

6.1 Datasets

The dataset consisted of 7,600 aggregated tweets from four corpora from the literature. These datasets were from different time frames and all come along with gold standards, which we used to perform quality checks and compute contest scores.

- **The Ritter corpus** by Reference [59], which consists of 2,400 tweets. The tweets were randomly sampled; however, the sampling method and original dataset size are unknown. It is estimated that the tweets were harvested around September 2010.
- **The Finin corpus** by Reference [26] consists of 441 tweets, which was the gold standard for a crowdsourcing annotation exercise. It is not stated how the corpus was created; however, our investigation puts the corpus between August to September 2008.

- **The MSM 2013 corpus**, the Making Sense of Microposts 2013 Concept Extraction Challenge dataset by Reference [10], which includes training, test, and gold data; for our experiments, we used the gold subset comprising 1,450 tweets.
- **The Wordsmith corpus**, reported in one of our previous works [20]. From the a corpus of six billion tweets, we sampled out 3,309 English ones using *reservoir sampling*—a family of randomized algorithms for sampling k items from a list S of n items.

6.2 Experimental Setup

6.2.1 Overview. Our experiments were designed for time-bound microtask crowdsourcing scenarios, which require speedy human-cognitive intervention. The workers were presented with a time-delayed view of streaming tweets. Each experiment competition ran for about 6min, during which, workers labelled as many tweets as possible, identifying instances of people, locations and organisations. The objective of each worker was to qualify for a set of top reward spots, which guaranteed a payoff. Unlike traditional paid microtasks, only a few of the workers got paid even if they produced correct annotations. This added to the urgency and the competitive drive that sped up the task to near real-time completion rates. We considered several questions, chief of which was: How does the reward spread affect the total task output of all the contest participants? Other questions naturally follow, such as the relationship between reward spread and annotation time, annotation quality, and worker exit. These formed the basis of our experimental conditions.

6.2.2 Timely Worker Recruitment Strategies. To achieve timely worker recruitment, we adopted a combination of several strategies:

- We posted our tasks repeatedly to maintain visibility within the recent tasks view of workers.
- We created multiple tasks that pointed to our Wordsmith platform, but ensured that workers could connect only once by keeping track of the connection IP address.
- We attempted to recruit, on the average, 10 times the number of workers than we required (around 1,000 workers).
- We posted the tasks in bits as a work around for the scheduling mechanism that Crowd-Flower uses in displaying unfinished tasks to new workers.
- We introduced an acoustic signal to notify workers once the requisite number of contestants had connected to the system.

Workers could see in real-time how many more contestants were required to connect before the task started. We also ensured that impatient workers were reconnected whenever they refreshed their screens; however, once the required number of workers were connected, no further connection was allowed.

6.2.3 Requester Parameters. In each experimental condition, the number of workers was 100, viewing a $k = 3$ neighbours leaderboard; the number of tasks (in our case Twitter posts) desired to be processed was 7,600; the window size was 200; the task unit time was 10s, the task unit size was 10. This means that a worker saw 10 tweets out of the current stream of 200 tweets for a period of 10s during which the worker was to label as many as possible before the next set of tasks arrived. The task arrival rate was 200/10, which means 20 tweets per second, and the total contest time was $(7,600 \times 10)/200 = 380s$ (6min, 20s).

Each contest had a payoff of \$0.10 for each prize payment, replicating [20, 23]. Our model rewarded each winner with the same payment, as opposed to other variants, which take into account the ranking of the participants or implement some other form of reward sharing. A reward spread of 1 stands for a winner-takes-all condition, in which only the top worker gets paid and with a

reward spread of 10, the top ten workers are paid. The parameters have been explained earlier in Section 5.2.

6.3 Experimental Design

6.3.1 Overview. We set up our experiment using the between-subjects design. Each worker could only be assigned to one of the experiment conditions. This was enforced using the worker's connection IP address as stated earlier in Section 6.2.2. Each user was randomly assigned to one of the experiment conditions and they were not able to pre-select which condition they ended in. The multiple tasks posted on CrowdFlower and referenced in Section 6.2.2 served as an identical entry point for each worker. From there, they are redirected to a centralized server that randomly assigns the workers to a certain treatment.

Our experiments had one independent variable: the reward spread, and one dependent variable: the task output. As we described in our theoretical modelling in Section 4, our hypothesis is that the reward spread can "explain" the changes in the contest output (referenced as the total annotations carried out by the workers). Therefore, the variable that a requester seeks to optimize is the total output, and they seek to control this by varying the reward spread.

- **Independent variable:** reward spread
- **Dependent variable:** task output

6.3.2 Conditions. The experiments in this article are compared against a control condition from an identical task reported in References [20, 23]. In that study, workers were also required to annotate named entities in tweets; however, they were not under any time constraint to carry out the task, and every worker that successfully completed the task got paid. In this article, however, the experimental conditions or treatments are based on the different reward spreads that crowd workers are randomly assigned to, with the contest setting inducing the time constraint.

- **Control Condition** Pay all workers
- **Experimental Condition 1** Pay top worker
- **Experimental Condition 2** Pay top 5 workers
- **Experimental Condition 3** Pay top 10 workers

We evaluated task performance in terms of task output (i.e., volume of work done). To assess the task output, we measured the total and average number of texts annotated by the annotators.

7 RESULTS

In this section, we present two sets of results. The first, from running simulations using the continuous-time Markov chains model developed in Section 4; and, the second sets of results from the experiment designs discussed in Sections 6.2 and 6.3.

7.1 Simulation Results

We ran a series of simulations to test the CTMC model for contest annotations tasks. The independent variable is the reward spread, which we hypothesise can explain the annotation rate of workers in the contest.

To aid visualisation, we ran the simulation using 5 annotators, 100 annotations, and 200 time steps. First, we generated λ_i values, which represent the behaviour of each of the 5 workers when they are within the payment or no payment leaderboard ranks.

Figure 5 displays randomly generated λ_i values for each worker $_i$. From Figure 2, we know λ_i influences τ_j^i , therefore, we can simulate a contest run to see how the worker's strategies pan out:

Fig. 5. λ rates for each worker representing how much effort a worker exerts when they are within and outside the reward spread.

Fig. 6. Contest simulation using generated λ rates.

worker₁ who exerts significantly more effort only when they are not getting paid, or worker₃ that exerts much effort even when they are in the lead.

The results in Figure 6 paint a picture of how the change in annotation rates within and outside the reward spread affect how the entire contest turns out. This becomes more evident as the race progresses, in this case, with worker₃ coming out at the top with their λ annotation strategy. The results also lend credence to mode of representing λ —for example, worker₂ who exerted *no* additional effort outside the reward spread (see Figure 5) came last in the contest run (see Figure 6). Therefore, we can extend this to model what would happen when *less* effort is exerted.

Therefore, Figure 6 illustrates that we can go from defining values for λ to simulating an entire contest run. The challenge then is to do the reverse, that is, to infer the values of λ from observing the transition τ s in a contest. Learning this would allow us to modify the reward spread in a real-life contest, which would in turn modify λ , resulting in a desired contest outcome. The desired outcome from a requester’s point of view would be to maximise the total annotations from all workers.

Fig. 7. Learned values of λ rates for each worker.

Using the loss function derived in Equation (2), we trained a model to learn the λ rates of each of our 5 workers. We used the worker's rank at each time step, the elapsed time, the number of annotations left, and whether the worker was within or outside the reward spread as features to train the model.

The results in Figure 7 illustrate that we can learn the λ rates of workers, which describe how they perform when they are within and outside the reward spreads. For example, observing the results of worker₀ in Figure 7 shows that the λ_0 rate when the worker is outside the reward spread is normally distributed around a mean value of 1.15, which is very close to the value of 1.12 generated and displayed in Figure 5. Similarly, the λ_0 rate when worker₀ is within the reward spread is normally distributed around 1.7, which is close to the generated value of 1.66.

7.2 Experiment Results

In the next sections, we discuss the results of our experiment runs, specifically, how the task output from contests yields a significant increase when compared to the baseline control condition. We then discuss the difference in annotation times across the control and experimental conditions, which is a direct consequence of the contest setting.

Then, we touch on the speed vs. quality trade-off that we incur as an aftereffect of the gain in task output. Finally, we revisit the subject of exits, which comes as a result of potential loss of utility on the part of the crowd workers. We tie this to the reward spread to explain how a change in reward spread can potentially stave off task abandonment.

7.3 Task Output

The results in Figure 8 illustrate how a change in reward spread leads to a change in the number of annotations submitted by workers. The central aim of the requester is to elicit as many annotations from the workers as possible (within the time constraint). Therefore, Figures 8(a) and 8(b) paint a picture of how an increase in the reward spread realises this objective. Both figures compare the total and distinct number of annotations in each of our experimental conditions, against the annotations from the baseline control condition.

(a) Total annotations by all workers

(b) Distinct annotations by all workers

(c) Average annotations per worker

(d) Number of active annotators

(e) Total annotations by top worker

(f) Total annotations by top 10 workers

Fig. 8. Task output.

From Figure 8(a), the null hypothesis would lead us to believe that the population mean of worker annotations (given by the control condition) is a total of 1,062 labels in the given time window. The alternative hypothesis from our experimental conditions suggest otherwise: that setting the task as a contest leads to a different mean, which varies by the reward spread. The resulting hypothesis that contests leads to more annotations is statistically significant. A corollary can be drawn from the distinct annotations in Figure 8(b).

We then carried out a one-way ANOVA test on the total annotations from the different experimental conditions to see if there was statistical significance as a result of a change in the reward spreads. Across the three conditions: pay top worker ($\mu = 1,322$, $\sigma = 100.40$), pay top 5 workers ($\mu = 1,502$, $\sigma = 132.22$), and pay top 10 workers ($\mu = 1,886$, $\sigma = 134.35$), we obtained results that give evidence to the significance of the results. The F-ratio value is 67.19, the p -value is < 0.00001 , and the result is significant at $p < 0.05$.

Figure 8(c) highlights the fairly constant average number of annotations by all the workers. This at first appears to contradict the results we described in Figures 8(a) and 8(b). However, Figure 8(d) illustrates that the number of active annotators varied across the three experimental conditions. This result also buttresses the significance of the different reward spreads in retaining more

Fig. 9. Average annotation time per entity.

annotators. Furthermore, a clearer picture of the number of annotations among workers is evidenced in Figures 8(e) and 8(f). The results make plain the effect of varying the reward spread, which has stronger impacts on the top workers.

7.4 Annotation Time

In Figure 9, we compare the time spent annotating an entity in the different experimental conditions against the time spent by workers in the control condition.

In the control condition, workers were not placed under any specific time constraint; however, inherent timely completion was required to receive their task compensation. Therefore, they still had an incentive to complete the annotations without any delay. In the control condition, workers needed on average 9.70s to identify and annotate an entity in a tweet. The calculated times in the experimental conditions were as follows: pay top worker ($\mu = 4.68s$, $\sigma = 0.08$), pay top 5 workers ($\mu = 4.65s$, $\sigma = 0.32$), and pay top 10 workers ($\mu = 4.77s$, $\sigma = 0.15$).

The results from the contests are significant when compared to the baseline condition, which did not have the explicit time constraint. Contest annotators needed an average of 4.70s to recognise a type of an entity while regular annotators spent 9.70s per entity. This is equivalent to an increased annotation factor of almost 2x across all contests. Contest annotators spent an average of 5.6s annotating one tweet (at 1.2 entities per tweet) and the top contestants averaged 2.69s per entity, resulting in a gain factor of 3.6x against the baseline.

7.5 Annotation Quality

An increase in annotation speed, however, led us to incur a hit in annotation quality as reflected in Figure 10.

The results presented in Figure 10 reflect a fine-grained analysis of precision, recall, and F1 scores on the four entity types (PER, ORG, LOC, and MISC) across the experimental and baseline conditions. The results follow the same pattern; i.e., the same entity types are equally difficult to identify across all conditions. However, the baseline condition, which was free from time constraints, resulted in better scores. The scores were computed against the gold standards published with the four datasets used for the experiments.

This result helps inform us as to what tasks contests might be better suited for. A requester seeking to adopt contests needs to be able to tolerate a certain level of false positives and false negatives, which would be compensated for by the timeliness of the responses. Verifying the quality of worker submissions (without using inter-annotator agreement scores), however, leads to a

Fig. 10. Annotation quality: Precision, Recall, and F1 scores for annotating Person (PER), Organisation (ORG), Location (LOC), and Miscellaneous (MISC) entities.

recursive chicken and egg scenario; for if we were able to automatically compute the quality of the annotations (i.e., by using a gold standard), then what did we need the crowd for? Or better yet, how can we design contests, and assign ranks in a way that creates quality output without falling prey to malicious workers who make quick but incorrect submissions?

In this case, we posit that the task design should be one that satisfies the following two conditions: (a) the worker solution could not have been easily computed automatically; but (b) the worker solution can be computationally verified easily without a gold standard. An example of such a scenario is in the fold.it game [12], where players fold the structures of selected proteins. Determining a protein's structure is computationally demanding; however, humans are able to intuitively arrive at the best structural composition. Once this has been done, the verification process becomes straightforward, making this task a prime candidate for carrying out this form of contest-based crowdsourcing without a gold standard.

7.6 Contest Exit Behaviour

Only a subset of contestants received a monetary payoff. The longer workers engaged with a task, the more utility they potentially lost given their ranking relative to the reward spread. Some workers therefore opted to leave the contest prematurely. Figure 11 presents the exit behaviour of workers across the three experimental conditions. The results imply that the beginning of the contest sees fewer exits with most of the contestants choosing to continue for up to 90% of the total time period. The exit rates increase toward the end. This behaviour would be peculiar to microtask contests unlike more longitudinal contests, such as presidential elections (studied by Reference [55]), in which most contenders exit at the beginning of the race leading up to much fewer participants at the final elections.

8 DISCUSSION

In this section, we highlight lessons learned from the experiments, the challenges of scaling contests to real-time annotations, and the limitations of our approach and results.

8.1 General Considerations

Crowdsourcing remains a mechanism with tremendous potential to grow economies by engaging a large workforce on demand and at scale. This translates into improved productivity and significant benefits for task requesters and crowd workers alike. Requesters gain task scalability, quick

Fig. 11. Number of active annotators at different stages of the contest. The stages are marked at 5% time interval steps throughout the contest.

completion, and lower price margins for their projects, while workers achieve additional income and learn new skills. In addressing crowdsourcing as a sociotechnical construct, in this article, we focus on one of the crowdsourcing challenges from the list discussed in Reference [37]: real-time crowd work. Tackling this challenge, we developed a model that encourages workers to perform well, even in time-sensitive situations, by using micropayments and contests.

8.2 Layering Incentives on Top of Payments

Our experiments showed that competitions could work in paid microtask crowdsourcing, both as an alternative to common task models and as a means to complete tasks faster. On average, workers carried out their tasks twice as fast as the baseline workers, while top performers were three times as fast. When we varied the number of workers who got paid in the experiments from 1 to 5 to 10, we noticed that the average volume of work also went up, with workers perceiving a greater chance to receive part of the available reward. As expected, in all the contest experiments, the gain in speed came at a cost and this was a trade-off for quality. When we compared the quality of submissions in the baseline experiments with that in the contest, we reported a drop in precision of entity identification and total recall of various entity types. In studying the exit patterns of workers from the task, we were keen to understand at what time in the contest do most players quit (potentially due to their inability to potentially finish in the reward spots). Given the intrinsic fun that comes from participation besides from receiving a potential financial payout, many workers (a minimum of 85% per condition) were willing to stay for almost the entire contest time. It was also quite intriguing that 49 workers rated the contest experience as 3.5 out of 5, despite only 9 of them receiving an actual payout, which makes us believe there must have been some form of satisfaction derived from the competition.

Our experiments give further insight into the dynamics of crowd behaviour. Workers posted on forums to notify others when a new contest was set up to quickly get 100 workers into the task. This was consistent to our prior image-labelling experiments that used Wordsmith—in Reference [20], we noted that some workers even promoted the game to their friends. During the study Wordsmith was used to label images by users who were not recruited via CrowdFlower who found out about the platform via social channels.

8.3 The Role of Financial Rewards

With the contests, we investigated the interplay between the reward spread and crowd behaviour. Allowing more workers to be eligible for payment improved the overall performance. With more winning spots came an increase in the total and average task output by all participants. Having a higher reward spread also ensured that workers stayed in the context and did not drop out early, thus leading to more tasks completed. This further meant that more effort was required to achieve one of the top spots in the ranking, which would receive a payment.

However, indefinitely increasing the reward spread would defeat the purpose of adopting a contest model, converging toward a traditional microtask system. Our results show that the linear result growth begins to break down as expected at some point. It is important to note that the motivation of crowd workers covers a wide spectrum of intrinsic and extrinsic factors, hence, having a wider reward spread (and a higher task threshold as seen in C6) led to a plummet in task output (over repeated experiment runs). An investigation into the discussion forums indicated that this experiment was probably less challenging—as stated by one worker, “... *to get into the top 10 is not too difficult...*”—and hence might not have been as attractive to top performers. Understanding where to draw the line would be the subject of further empirical studies paired up with theoretical analysis.

8.4 Scaling Toward Real-time Annotations

Microtask crowdsourcing has often been praised for its ability to produce results quickly and accurately. Yet, an increasing number of applications come in with much harder time constraints, which push the boundaries of the traditional microtask model to deliver in seconds or less [4]. Examples of such real-time crowdsourcing applications include machine learning for image recognition [69] and text-to-speech conversion [44]; accessibility design [7, 43, 44]; and disaster management [31].

The study lends support to our hypothesis that the microtask contest model leads to timely task completion without the associated overhead costs. Our results show an increase in speed by an average factor of 2x, and up to 3x more among the top performers. These metrics could be used to inform decisions on the number of workers that would be needed to annotate the entire dataset of 7,600 tweets in real-time—taking into consideration the total number of non-unique annotations and worker exit e.g., in experiment C5, workers annotated 1981 tweets, equivalent to 384 workers required to make one pass at the entire stream. Our experiments were carried out on a stream of 20 tweets per second; the live Twitter stream is currently estimated at 6,000 tweets per second or 2,400 English tweets per second (40% of the full stream) [24]. Annotating 10% of the live stream could be potentially carried out by designing a contest for 4,600 workers (i.e., $384 \text{ workers} \times ((10\% \text{ of } 2,400)/20 \text{ tweets per second}))$).

8.5 A Diverse Crowd

Across all experimental conditions, we have observed a range of responses to the incentives we layered on top of micropayments, which varied depending on a range of factors, including the amount of work already completed and the likelihood to be rewarded. During the contests, top contributors showed a different behaviour in the conditions with larger reward spreads. These findings suggest the need for additional experiments, which would investigate the effects of the different incentives on different groups of workers—for example, those who contributed a lot in terms of output but also those who over achieve in terms of quality. This would potentially allow us to predict exit behaviour more accurately, and offer finer grained incentives—for instance, one could imagine different bonus levels based on how many objects a worker has annotated so far.

8.6 When Is This the Right Model?

Our experiments and discussions have brought us to a number of observations: Layering contest-based incentives on top of financial payments can help unlock the potentials of real-time annotations. However, interpreting the trade-off between quality, speed, and task output is best understood on a case-by-case basis. We describe two criteria that can help requesters decide if leveraging contests is the right model for their task. The criteria are: price sensitivity and trade-off suitability. We further touch upon price sensitivity, payments, and ethical considerations in Section 8.8; however, this criterion suggests that requesters use contests for tasks that have a high intrinsic appeal. Examples of these would be large-scale volunteer efforts (such as annotations after the Haiti Earthquake [75]), or platforms with existing contest models (such as TopCoder [1]).

With regards to trade-off suitability (between quality, speed and task output), we describe two scenarios where this model could find application. (i) Tasks with automatic quality verification (e.g., in the protein folding game fold.it [12]). In these cases, speed and task output become directly correlated and the requester immediately benefits from deploying a contest model without the quality constraints. (ii) Tasks where quality can be compensated by larger volumes and timely responses (e.g., disaster response annotations [2, 31] for earthquakes, floods, etc). In these tasks, strategies such as majority voting can be adopted to further identify high quality responses. In both cases, our recommendation would be that the requester select the largest reward spread based on their available budget and the prevailing payment for the task.

8.7 Application to Active Learning

The contest model also finds potential application in active learning. Active learning is a machine-learning paradigm that seeks to minimise the time and crowdsourced labels required to train a machine-learning model [64]. Since the objective of active learning is to scale down the the amount of labels required, it is of necessity that the labels be of high quality. However, in large-scale real-time active learning and crowdsourcing scenarios [19, 68, 72], a single “oracle” providing labels will not suffice. Just as active learning algorithms refine the model and suggest the few training examples required, crowdsourcing contests refine the annotator pool to provide the “expert crowd” required to improve the model’s decision boundary. Applying the contest model in an active learning setting suggests that contest incentives do not always have to optimise for speed (unlike in the real-time scenario). This opens up avenues for future research into incentives in high-value judgement tasks that are aligned more toward annotator skill and less toward completion speed.

8.8 Payments and Ethical Considerations

Crowd worker motivation remains a constant research area in understanding why people partake in microtasks. This is essential to design systems that are fair and rewarding to the workers. A requester would always seek to minimise cost, however, a worker’s complete range of motivations is not yet fully understood. Our investigations revealed that the task model was relatively well received: in one of the baseline experiments, 87 workers rated the payment of \$0.10 as 4 out of 5, while in one of our contests, 49 workers rated it as 3.5 on the same scale despite only 9 of them receiving an actual payout. (The numbers 87 and 49 represent all the workers who gave feedback on the task.) Furthermore, workers seemed to be eager to return—as stated in a crowd discussion forum, one of the participants posted *“Hit the top 10 today. I will hunt this problem again,”* while another one felt let down when they couldn’t be among the 100 contestants: *“Yesterday I came across this, but [they] recruited 100 people, [I] was not allowed to play.”*

On task payment, we favoured a higher than average payment of \$0.10 as against the annotation averages reported by Reference [18]; however, it would be interesting to see the effects of increased

payments on the results. In a related set of experiments, we noted that raising the reward to as much as \$0.25 rather created an anchoring effect than a trend in results. We would like to investigate this further and also analyse the effect of task payment ordering. A greater understanding of worker intrinsic motivations would help in the design of better payment schemes, wrapped around task models that workers find inherently engaging and rewarding; and lend insights into the ongoing debate on ethical and fair crowdsourcing [34].

8.9 Limitations

While the results of our experiments were very promising, we are aware of several limitations of our work. The experiments are run on text annotation tasks. It would be interesting to test the models on the other five types of tasks from Reference [30]. While our experiments and theories could be easily adapted to similar task types, in particular, to output-agreement ones [69], they are less suitable for settings in which a diversity of answers is sought. As a result, there is much latitude for future work across the different incentive schemes and the concept of furtherance incentives as a whole. The tasks we tested do not require a high cognitive load, time investment or creativity. Other types of tasks might be less suitable for a casual game model like the one in Wordsmith. Moreover, the results will probably vary when the tasks are carried out by a community who is intrinsically motivated (for example, to advance science as in citizen science projects, or to help, as in disaster relief). Those motivations would probably overshadow the simpler mechanisms tested here. In addition, such intrinsic motivation settings have been shown to be more prone to overjustification effects, and any attempt to use incentives to reward contributions might have negative effects.

Just like most empirical research in paid microtask crowdsourcing, none of the experiments considered temporal effects. While contests may produce positive effects short-term, it would be interesting to understand whether workers would continue to engage more with no additional incentives as tasks run over longer periods of time. One common concern related to virtual rewards is that they are not always aligned with the needs and values of the user and that this mismatch becomes more obvious with time [17]. In our experiments, we have not carried out an in-depth study of the types of workers who used Wordsmith or try to understand what drives them. While the experiments were successful, we anticipate that we would need to have a much more nuanced approach to contest-based incentives when workers engage with us over longer periods of time.

Our contest experiments allowed us to get a sense of the number of contributors required to annotate a substantial share of the data feed in near-real-time. However, our findings might not transfer to scenarios with contests that stretch over days or longer. In the same time, in a real-world scenario, most human labelling projects that operate with large amounts of data would also include an algorithmic component, trained by labels created by the crowd.

Finally, our work assumes that there is a gold standard to verify crowd answers. This is particularly important in time-sensitive scenarios.

9 CONCLUSIONS AND FUTURE WORK

Our experiments have explored the effects of layered incentive mechanisms on a base financial payment model on crowd performance and engagement. The results suggest that we can build better crowdsourcing systems for task requesters and workers alike; and that money need not be the only currency for transacting on paid microtask platforms

In line with the research on timely task completion, we intend to carry out new real-time crowdsourcing experiments that feature automatically verifiable task output, which is currently a limitation of our current approach. We would also carry out cross studies between social incentives-based microtasks and contest-based tasks in a bid to understand how the former could be used

to facilitate the latter. Our approach to use competitions to deliver timely annotations naturally engenders worker exit given the varying utilities gained by individual workers during the contest. By doing so, it offers us a way to gain insight into how the different parameters we use to our probabilistic exit predictor impact behaviour. This can be further explored by using the experiment data to learn suitable parameters for our CTMC model.

Overall, our results reinforced the role of motivation and incentives beyond cash payments as the primary means to reward participation on some of the most popular crowdsourcing platforms of our times. Our experiments give crowdsourcing researchers and platform designers new empirical evidence into crowdsourcing as a fully fledged sociotechnical construct with challenges requiring as much a social touch as the technical. Worker motivations cover a wide spectrum, which extends beyond the simplistic spectrum of “fun and money.” Understanding how to keep the crowd engaged is essential not just to yield superior returns for requesters but also to create a more humane experience, as crowdsourcing continues to take its place as the model for the future of work.

ACKNOWLEDGMENT

We acknowledge Karen Hovsepien for all the helpful conversations.

REFERENCES

- [1] Nikolay Archak. 2010. Money, glory, and cheap talk: Analyzing strategic behavior of contestants in simultaneous crowdsourcing contests on TopCoder.com. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 21–30.
- [2] Zahra Ashktorab, Christopher Brown, Manojit Nandi, and Aron Culotta. 2014. Tweedr: Mining Twitter to inform disaster response. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM'14)*.
- [3] James Bennett and Stan Lanning. 2007. The Netflix prize. In *Proceedings of the Knowledge Discovery and Data Mining Cup and Workshop (KDD'07)*. 35.
- [4] Michael S. Bernstein, Joel Brandt, Robert C. Miller, and David R. Karger. 2011. Crowds in two seconds: Enabling real-time crowd-powered interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. ACM, 33–42.
- [5] Michael S. Bernstein, David R. Karger, Robert C. Miller, and Joel Brandt. 2012. Analytic methods for optimizing realtime crowdsourcing. *Collect. Intell.* (2012). <http://www.ci2012.org/>.
- [6] Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2015. Soylent: A word processor with a crowd inside. *Commun. ACM* 58, 8 (2015), 85–94.
- [7] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, et al. 2010. VizWiz: Nearly real-time answers to visual questions. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*. ACM, 333–342.
- [8] Kevin J. Boudreau, Nicola Lacetera, and Karim R. Lakhani. 2011. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Manage. Sci.* 57, 5 (2011), 843–863.
- [9] Randy Burkett. 2013. An Alternative Framework for Agent Recruitment: From MICE to RASCLS. Retrieved from <https://calhoun.nps.edu/handle/10945/43831>.
- [10] Amparo Elizabeth Cano Basave, Andrea Varga, Matthew Rowe, Milan Stankovic, and Aba-Sah Dadzie. 2013. Making sense of microposts (# MSM2013) concept extraction challenge. In *Proceedings of the Concept Extraction Challenge at the Workshop on 'Making Sense of Microposts' Co-Located with the 22nd International World Wide Web Conference (WWW'13)*, Amparo E. Cano, Matthew Rowe, Milan Stankovic, and Aba-Sah Dadzie (Eds.). CEUR-WS.org, 1–15.
- [11] Shuchi Chawla, Jason D. Hartline, and Balasubramanian Sivan. 2015. Optimal crowdsourcing contests. *Games Econ. Behav.* 113 (2015), 80–96. <https://doi.org/10.1016/j.geb.2015.09.001>
- [12] S. Cooper, F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popović, et al. 2010. Predicting protein structures with a multiplayer online game. *Nature* 466, 7307 (2010), 756–760.
- [13] Peng Dai, Jeffrey M. Rzeszotarski, Praveen Paritosh, and Ed H. Chi. 2015. And now for something completely different: Improving crowdsourcing workflows with micro-diversions. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 628–638.
- [14] Emmanuel Dechenaux, Dan Kovenock, and Roman M. Sheremeta. 2015. A survey of experimental research on contests, all-pay auctions and tournaments. *Exper. Econ.* 18, 4 (2015), 609–669.

- [15] Edward L. Deci and Richard M. Ryan. 2010. *Self-Determination*. John Wiley & Sons.
- [16] Gianluca Demartini, Djellel Eddine Difallah, and Philippe Cudré-Mauroux. 2013. Large-scale linked data integration using probabilistic reasoning and crowdsourcing. *Very Large Data Base J.* 22, 5 (2013), 665–687.
- [17] Sebastian Deterding, Dan Dixon, Rilla Khaled, and Lennart Nacke. 2011. From game design elements to gamefulness: Defining gamification. In *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments*. ACM, 9–15.
- [18] Djellel Eddine Difallah, Michele Catasta, Gianluca Demartini, Panagiotis G. Ipeirotis, and Philippe Cudré-Mauroux. 2015. The dynamics of micro-task crowdsourcing: The case of Amazon MTurk. In *Proceedings of the 24th International Conference on World Wide Web*. 238–247.
- [19] Oluwaseyi Feyisetan, Thomas Drake, Borja Balle, and Tom Diethe. 2019. Privacy-preserving active learning on sensitive data for user intent classification. *arXiv preprint arXiv:1903.11112*.
- [20] Oluwaseyi Feyisetan, Markus Luczak-Rösch, Elena Simperl, Ramine Tinati, and Nigel Shadbolt. 2015. Towards hybrid NER: A study of content and crowdsourcing-related performance factors. In *The Semantic Web: Latest Advances and New Domains*. Springer, 525–540.
- [21] Oluwaseyi Feyisetan and Elena Simperl. 2016. Please stay vs let’s play: Social pressure incentives in paid collaborative crowdsourcing. In *Proceedings of the International Conference on Web Engineering*. Springer, 405–412.
- [22] Oluwaseyi Feyisetan and Elena Simperl. 2017. Social incentives in paid collaborative crowdsourcing. *ACM Trans. Intell. Syst. Technol.* 8, 6 (2017), Article no 73.
- [23] Oluwaseyi Feyisetan, Elena Simperl, Markus Luczak-Roesch, Ramine Tinati, and Nigel Shadbolt. [n.d.]. An extended study of content and crowdsourcing-related performance factors in named entity annotation. *Semantic Web* 1–25.
- [24] Oluwaseyi Feyisetan, Elena Simperl, Ramine Tinati, Markus Luczak-Roesch, and Nigel Shadbolt. 2014. Quick-and-clean extraction of linked data entities from microblogs. In *Proceedings of the 10th International Conference on Semantic Systems (SEM’14)*. ACM, 5–12.
- [25] Oluwaseyi Feyisetan, Elena Simperl, Max Van Kleek, and Nigel Shadbolt. 2015. Improving paid microtasks through gamification and adaptive furtherance incentives. In *Proceedings of the 24th International Conference on World Wide Web*. 333–343.
- [26] Tim Finin, Will Murnane, Anand Karandikar, Nicholas Keller, Justin Martineau, and Mark Dredze. 2010. Annotating named entities in Twitter data with crowdsourcing. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*. Association for Computational Linguistics, 80–88.
- [27] Brent Frei. 2009. Paid crowdsourcing. *Current State & Progress toward Mainstream Business Use, Smartsheet.com Report, Smartsheet.com* 9 (2009).
- [28] Bruno S. Frey and Reto Jegen. 2001. Motivation crowding theory. *J. Econ. Surveys* 15, 5 (2001), 589–611.
- [29] Drew Fudenberg and Jean Tirole. 1986. A theory of exit in duopoly. *Econometr.: J. Econometr. Soc.* (1986), 943–960.
- [30] Ujwal Gadiraju, Ricardo Kawase, and Stefan Dietze. 2014. A taxonomy of microtasks on the web. In *Proceedings of the 25th ACM Conference on Hypertext and Social Media*. ACM, 218–223.
- [31] Huiji Gao, Geoffrey Barbier, and Rebecca Goolsby. 2011. Harnessing the crowdsourcing power of social media for disaster relief. *IEEE Intell. Syst.* 3 (2011), 10–14.
- [32] Benjamin M. Good and Andrew I. Su. 2011. Games with a scientific purpose. *Genome Biol.* 12, 12 (2011), 1.
- [33] Jeff Howe. 2006. The rise of crowdsourcing. *Wired Mag.* 14, 6 (2006), 1–4.
- [34] Lilly C. Irani and M. Silberman. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 611–620.
- [35] Nicolas Kaufmann, Thimo Schulze, and Daniel Veit. 2011. More than fun and money. Worker motivation in crowdsourcing-A study on mechanical turk. In *Proceedings of the Americas Conference on Information Systems (AMCIS’11)*, Vol. 11. 1–11.
- [36] Aniket Kittur. 2010. Crowdsourcing, collaboration and creativity. *ACM Crossroads* 17, 2 (2010), 22–26.
- [37] Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. ACM, 1301–1318.
- [38] Aniket Kittur, Boris Smus, Susheel Khamkar, and Robert E. Kraut. 2011. Crowdforge: Crowdsourcing complex work. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. ACM, 43–52.
- [39] Markus Krause and René Kizilcec. 2015. To play or not to play: Interactions between response quality and task complexity in games and paid crowdsourcing. In *Proceedings of the 3rd AAAI Conference on Human Computation and Crowdsourcing*.
- [40] Robert E. Kraut, Paul Resnick, Sara Kiesler, Moira Burke, Yan Chen, Niki Kittur, Joseph Konstan, Yuqing Ren, and John Riedl. 2012. *Building Successful Online Communities: Evidence-based Social Design*. MIT Press.
- [41] Vijay Krishna and John Morgan. 1997. An analysis of the war of attrition and the all-pay auction. *J. Econ. Theory* 72, 2 (1997), 343–362.

- [42] Walter S. Lasecki, Christopher Homan, and Jeffrey P. Bigham. 2014. Architecting real-time crowd-powered systems. *Hum. Comput.* 1, 1 (2014).
- [43] Walter S. Lasecki, Christopher Miller, Adam Sadilek, Andrew Abumoussa, Donato Borrello, Raja Kushalnagar, and Jeffrey Bigham. 2012. Real-time captioning by groups of non-experts. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*. ACM, 23–34.
- [44] Walter S. Lasecki, Christopher D. Miller, and Jeffrey P. Bigham. 2013. Warping time for more effective real-time crowdsourcing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2033–2036.
- [45] Walter S. Lasecki, Kyle I. Murray, Samuel White, Robert C. Miller, and Jeffrey P. Bigham. 2011. Real-time crowd control of existing interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. ACM, 23–32.
- [46] Edith Law, Ming Yin, Kevin Chen Joslin Goh, Michael Terry, and Krzysztof Z. Gajos. 2016. Curiosity killed the cat, but makes crowdwork better. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4098–4110.
- [47] Tracy Xiao Liu, Jiang Yang, Lada A. Adamic, and Yan Chen. 2011. Crowdsourcing with all-pay auctions: A field experiment on Taskcn. *Proc. Assoc. Info. Sci. Technol.* 48, 1 (2011), 1–4.
- [48] Thomas W. Malone, Robert Laubacher, and Chrysanthos Dellarocas. 2010. The collective intelligence genome. *MIT Sloan Manage. Rev.* 51, 3 (2010), 21.
- [49] David Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a turker. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 224–235.
- [50] Winter Mason and Duncan J. Watts. 2010. Financial incentives and the performance of crowds. *ACM SigKDD Explor. Newslett.* 11, 2 (2010), 100–108.
- [51] Elisa D. Mekler, Florian Brühlmann, Klaus Opwis, and Alexandre N. Tüch. 2013. Disassembling gamification: The effects of points and meaning on user motivation and performance. In *Proceedings of the Conference on Human Factors in Computing Systems—Extended Abstracts (CHI'13)*. ACM, 1137–1142.
- [52] Pietro Michelucci. 2016. *Handbook of Human Computation*. Springer.
- [53] Benny Moldovanu and Aner Sela. 2001. The optimal allocation of prizes in contests. *Amer. Econ. Rev.* (2001), 542–558.
- [54] Benny Moldovanu, Aner Sela, and Xianwen Shi. 2012. Carrots and sticks: Prizes and punishments in contests. *Econ. Inquiry* 50, 2 (2012), 453–462.
- [55] B. Norrander. 2006. The attrition game: Initial resources, initial contests and the exit of candidates during the U.S. presidential primary season. *Brit. J. Polit. Sci.* 36, 3 (2006), 487–507.
- [56] Oded Nov, Ofer Arazy, and David Anderson. 2014. Scientists@ home: What drives the quantity and quality of online citizen science participation? *PloS One* 9, 4 (2014), e90375.
- [57] Jordan Raddick, Georgia Bracey, Pamela L. Gay, Chris J. Lintott, Phil Murray, Kevin Schawinski, Alexander S. Szalay, and Jan Vandenberg. 2010. Galaxy zoo: Exploring the motivations of citizen science volunteers. *Astron. Edu. Rev.* 9, 1 (2010).
- [58] Neal Reeves, Peter West, and Elena Simperl. 2018. “A game without competition is hardly a game”: The impact of competitions on player activity in a human computation game. In *Proceedings of the 6th AAAI Conference on Human Computation and Crowdsourcing (HCOMP'18)*. AAAI.
- [59] Alan Ritter, Sam Clark, Oren Etzioni, et al. 2011. Named entity recognition in tweets: An experimental study. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1524–1534.
- [60] Markus Rokicki, Sergiu Chelaru, Sergej Zerr, and Stefan Siersdorfer. 2014. Competitive game designs for improving the cost effectiveness of crowdsourcing. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 1469–1478.
- [61] Markus Rokicki, Sergej Zerr, and Stefan Siersdorfer. 2015. Groupsourcing: Team competition designs for crowdsourcing. In *Proceedings of the 24th International Conference on World Wide Web*. 906–915.
- [62] Razvan Rughinis. 2013. Gamification for productive interaction: Reading and working with the gamification debate in education. In *Proceedings of the 8th Iberian Conference on Information Systems and Technologies (CISTI'13)*. IEEE, 1–5.
- [63] Katie Seaborn and Deborah I. Fels. 2015. Gamification in theory and action: A survey. *Int. J. Hum.-Comput. Studies* 74 (2015), 14–31.
- [64] Burr Settles. 2012. Active learning. *Synth. Lect. Artif. Intell. Mach. Learn.* 6, 1 (2012), 1–114.
- [65] Victor S. Sheng, Foster Provost, and Panagiotis G. Ipeirotis. 2008. Get another label? Improving data quality and data mining using multiple, noisy labelers. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 614–622.
- [66] John C. Tang, Manuel Cebrian, Nicklaus A. Giacobe, Hyun-Woo Kim, Taemie Kim, and Douglas Beaker Wickert. 2011. Reflecting on the DARPA red balloon challenge. *Commun. ACM* 54, 4 (2011), 78–85.

- [67] Ramine Tinati, Markus Luczak-Roesch, Elena Simperl, and Wendy Hall. 2017. An investigation of player motivations in Eyewire, a gamified citizen science project. *Comput. Hum. Behav.* 73 (2017), 527–540.
- [68] Sudheendra Vijayanarasimhan and Kristen Grauman. 2014. Large-scale live active learning: Training object detectors with crawled data and crowds. *Int. J. Comput. Vis.* 108, 1–2 (2014), 97–114.
- [69] Luis von Ahn and Laura Dabbish. 2004. Labeling images with a computer game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'04)*. ACM, New York, NY, 319–326.
- [70] Luis Von Ahn and Laura Dabbish. 2008. Designing games with a purpose. *Commun. ACM* 51, 8 (2008), 58–67.
- [71] Xianzhong Xie, Helin Yang, Athanasios V. Vasilakos, and Lu He. 2014. Fair power control using game theory with pricing scheme in cognitive radio networks. *J. Commun. Netw.* 16, 2 (2014), 183–192.
- [72] Yan Yan, Romer Rosales, Glenn Fung, and Jennifer G. Dy. 2011. Active learning from crowds. In *Proceedings of the International Conference on Machine Learning (ICML'11)*, Vol. 11. 1161–1168.
- [73] Dejun Yang, Xi Fang, and Guoliang Xue. 2012. Game theory in cooperative communications. *IEEE Wireless Commun.* 19, 2 (2012).
- [74] Gabe Zichermann. 2011. Gamification has issues, but they aren't the ones everyone focuses on. O'ReillyRadar.
- [75] Matthew Zook, Mark Graham, Taylor Shelton, and Sean Gorman. 2010. Volunteered geographic information and crowdsourcing disaster relief: A case study of the Haitian earthquake. *World Med. Health Policy* 2, 2 (2010), 7–33.

Received February 2018; revised March 2019; accepted March 2019