TAKING REJECTION PERSONALLY: AN ETHICAL ANALYSIS OF WORK REJECTION ON AMAZON MECHANICAL TURK

Research paper

Johnstone, David, Victoria University of Wellington, New Zealand, david.johnstone@vuw.ac.nz

Tate, Mary, Queensland University of Technology, Australia, mary.tate@qut.edu.au

Fielt, Erwin , Queensland University of Technology, Australia, erwin.fielt@qut.edu.au

Abstract

Amazon Mechanical Turk (AMT) is a popular microtasking platform connecting those who need work done with those willing to do it. However AMT has come under increasing scrutiny for the way workers are treated on its platform. This paper examines one particular component of AMT's microtasking process by conducting a formal ethical analysis, using Tavani's Comprehensive Cyberethics Methodology, on the way work submitted for evaluation is assessed and either accepted or rejected. The study finds that the system is skewed in favour of those evaluating the work, with little recourse for workers, and an apparent disinterest from AMT. This paper contributes to the continuing debate over the governance of, and responsibilities for, those engaged in digital work through such platforms.

Keywords: Crowdsourcing, microtasking, Amazon Mechanical Turk, ethics.

1   Introduction

The term *crowdsourcing* was originally coined by Howe (2006) in an article in Wired magazine, and the ensuing plethora of definitions generally agree on the key features: setting a work task through an open call, using an Internet-based platform, for an online community (the 'crowd') to perform the task. The platform is the medium enabling communication between the caller (who sets the work task) and the crowd (who agree to work on the task and deliver the results back to the caller).

Different types of crowdsourcing have evolved, such as contests, collaborative communities, complementors, and labour markets (Boudreau and Lakhani, 2013). According to Boudreau and Lakhani, labour markets match buyers and sellers of services and employ conventional contracting for services rendered. They perform on-demand matching to give immediate support at an unprecedented scale depending heavily on sophisticated technology infrastructure and platform design to govern transaction effectively and efficiently. Due to their high automation and low transaction cost, labour markets allow for *microtasking* as one specific type of crowdsourcing, where small pieces of work that are not readily automated (microtasks) are distributed to individuals in a crowd. Microtasks are typically short in duration, repetitive, and offer low rates of pay. Types of microtasks vary considerably, but popular options include tagging photographs, translating pieces of text, and completing short surveys. Those requesting work will pay those completing the work, and the platform provider, who brings them together as a service provider within an online work marketplace, is paid a percentage.

The largest, and arguably best known, online platform within the microtask marketplace is *Amazon Mechanical Turk* (AMT). The business model behind AMT has barely changed since its inception.

However, the number of people and organisations offering work (the *Requesters*) and the number of individuals willing to engage in work on AMT have continued to increase (Pittman and Sheehan, 2016). The benefits for the Requesters centre on the access to a labour workforce that is scalable and available on-demand, incurring minimal transaction costs and logistical hurdles (Felstiner, 2011). Workers have stated they enjoy the option of working anywhere, anytime on a type of work of their own choosing – with a computer and Internet connection the only prerequisites (Felstiner, 2011).

There has been controversy around microtasking, with frequent accusations of exploitation, particularly focused on the low pay that workers receive, in both the popular media (e.g. Harris, 2014; Kavanaugh, 2017) and academia (e.g. Ettlinger, 2016; Pittman and Sheehan, 2016; Williamson, 2016). By claiming Workers are contractors, and not employees, legal requirements around pay and working conditions are largely avoided by both the Requesters and the platform owners. Although AMT has been under the most scrutiny in this respect, it was a different platform owner offering similar services, Crowdflower, that was eventually taken to court over the employment issue[1].

However, it has been reported that there are other controversial practices involved in the AMT process around microtasking (Silberman and Irani, 2016). One, in particular, involves the Requester's ability to reject the work submitted by individual workers, with a corresponding loss of the Worker's pay, time and reputation. This study examines this practice in depth, performing an ethical analysis of the roles different stakeholders play. The aims of the paper are to establish how ethical analysis can inform practices for managing acceptance/rejection of micro-tasking-based crowdwork, identify situations that give rise to ethical issues, and discuss the options available to ameliorate situations that are determined to be unfair.

The remainder of this paper is structured as follows. The AMT process for microtasking is outlined in detail, including the way microtask work is evaluated and either accepted or rejected. Following this, we draw on past work and explain the different reasons why microtask work might be rejected, and we organise these reasons into a classification scheme. We then employ a systematic analysis of these classes of work rejection using Tavani's (2013) Comprehensive Cyberethics Methodology (CCM). The results are outlined on a set-by-step basis, and their implications discussed. We finish with a set of conclusions.

2 The AMT Process

AMT is a platform acting as a work allocation marketplace. People seeking microtasking work register with AMT, providing information about themselves, including demographic information and their skillset. A Requester, who may be an individual, organisation, or representative agency, will register with AMT as someone offering microtasking work, though they are not required to provide as much information about themselves as prospective workers do. The Requester posts a description of the work to be done (the *human intelligence task*, or HIT) on the AMT platform, clearly specifying (ideally) what has to be done, what level of quality is required, and what will be paid.

The HIT is made available to prospective workers, who decide whether or not to agree to perform the work. The Requester may specify the HIT is only to be seen by workers who satisfy certain criteria (for example, an academic survey HIT may be restricted to certain demographics). AMT also maintains a Worker *reputation* measure (based on how many pieces of their work were accepted by the Requesters in the past), which is promoted by AMT as a measure of worker quality – another feature Requesters can use to limit HIT availability. AMT awards a special status of 'Master' for those Workers who, over time, attain sufficient success with accepted work, though Amazon will not reveal

---

[1] This case was a class action representing a group of some 20,000 Crowdflower workers, arguing that Crowdflower were in breach of the US Fair Labor Standards Act, and was settled in the plaintiffs' favour. Details can be found in: https://www.overtimepaylaws.org/federal-court-approves-settlement-in-crowdsourcing-labor-company-wage-suit/

how, exactly, this status is achieved. Requesters can ask AMT to promote their tasks to only those with good reputations (especially Masters), though a higher rate of pay is usually required.

Once a Worker has accepted a HIT, they have to complete the work and submit it to the Requester (via the AMT platform) within the specified timeframe. The Requester is then in a position to evaluate the quality of the work and decide whether to accept it or reject it. If the work is accepted, the Requester may use the work for their own purposes, and the Worker will automatically be paid using funds supplied by the Requester to AMT at the time the HIT was first posted. This would seem to complete the contractual obligations of the Requester/Worker relationship.

However, as specified in the *AMT Participation Agreement[2]*, Requesters also have the right to reject any work submitted if they are "not reasonably satisfied" with the quality of the work submitted. Clearly the intention is to give Requesters the right to refuse payment for work that is seen to be substandard. In practice, it turns out there are many reasons why work might be rejected, some of which may, or may not, be argued to be unfair. The focus of this paper is on some of the ethical issues that may arise when submitted work is rejected.

AMT is paid a commission by the Requester (at the time of writing, usually 20% of the payments to Workers, though this can vary), but only for work that has been accepted. It should also be noted that a more informal reputation system about Requesters is provided by software, referred to as Turkopticon, created by Stanford researchers Silberman and Irani (2016) for AMT Workers to use, in an attempt to redress the balance in reputation tracking.

3 Why is the work rejected?

There are many possible reasons why work submitted by a Worker is rejected by the Requester, and these have been comprehensively listed by those who have explored AMT work practices (Silberman and Irani, 2016), AMT and employment law (Felstiner, 2011) and managing the risks around work rejection (McInnis et al., 2016). From these sources, six explicit types of work rejection have been identified, and we have grouped these into three classes. These are described as follows:

A. The Worker did not provide quality work. There are two possible reasons for this:

    i. The Worker may have tried to meet expectations, but either did not have the capacity (e.g. time, knowledge, experience, effort, skills, working environment), or simply did not end up putting in sufficient time or effort, to do so.

    ii. The Worker did not even try to meet expectations. In this case it usually means the Worker sought to deceive the Requester, for example by submitting random responses to survey questions, quickly, in order to make more money in a given time period, hoping the Requester would employ an inadequate evaluation system to assess the quality of responses.

  B. The Requester did not provide an adequate opportunity for the Worker to meet expectations. In this case, both the Requester and the Worker may have had good intentions, but the design and/or presentation of the HIT meant the Worker was not able to meet expectations. This can occur when, for example, the Requester is relatively inexperienced. More specifically:

    iii. The Requester may provide a task description that is vague or ambiguous. If the instructions and/or evaluation criteria are not clear, it makes it harder for the Worker to determine how best to meet the Requester's expectations.

    iv. The Worker may determine that there is simply too much to do to complete the task, either because it cannot be justified in terms of the rate of pay, or cannot be reasona-

---

[2] https://www.mturk.com/mturk/conditionsofuse

bly finished within the timeframe set by the Requester. As a result, the Worker may submit work that is incomplete or rushed, or even abandon the attempt (which has the same effect as a rejection).

C.   The Worker provided quality work, but the Requester rejected it anyway. Possible reasons include:

v.   The Requester may have employed poorly designed and/or written software, using the Application Programming Interface (API) supplied by AMT, that automates the evaluation of Worker responses that might, on at least some occasions, reject work that should not have been rejected. Automated evaluation techniques include, for example, (a) gold standard questions - tasks within a HIT that Requesters already know the correct answer for (e.g. in a survey), thus establishing a check that Workers are providing correct answers, or (b) verification systems such as assigning the same task (e.g. tagging a photograph) to two workers and comparing the answers (if they are different, a third worker is assigned the task).

vi.   The Requester may deliberately reject submissions to avoid paying for work they intend to make use of anyway.

In summary, Type A situations suggest it is fair that the work is rejected. Type C situations suggest it is clearly unfair that the work is rejected. Type B situations also suggest it is unfair that the work is rejected, but the degree of unfairness may be difficult to establish if there is any disagreement between the Requester and the Worker about the task as described and presented.

It should be noted, at this point, that there is a third status that applies to submitted work when it is neither accepted nor rejected. This occurs when the Requester does not make a decision on the submission. In these situations, AMT will eventually approve the HITS and pay the Workers. The time elapsed until this occurs (known as 'automatic approval' or 'AA' time) defaults to 30 days, though Requesters may set AA time to be less if they wish when posting the HIT.

In this study we ask the following questions. First, what are the ethical implications of these different situations? Second, what ethical responsibilities do the three stakeholders involved (the Worker, the Requester, and AMT) have to address any ethical implications?

4   Method

We investigate the rejection of Worker submissions through the application of Tavani's (2013) Comprehensive Cyberethics Methodology (CCM). Tavani (2013) defines *cyberethics* "as the study of moral, legal and social issues involving cybertechnology" (p.4), where *cybertechnology* "refers to a wide range of computing and communication devices, from stand-alone computers, to connected, or networked, computing and communication technologies" (p.4). CCM is particularly well-suited to this study as it adapts a 'standard model used in applied ethics' to include additional characteristics of cybertechnologies that are not considered or emphasised in the standard model, such as policy vacuums and conceptual muddles and the moral implications of features embedded within the technology. The description of the CCM process that follows explains these concepts.

The purpose of CCM is to provide a method for guidance "in the identification, analysis, and deliberation processes" (p.27) required to evaluate an ethical issue and determine possible action. The following describes and explains the steps involved in CCM, based on Tavani's material.

Step 1: Identify a practice involving cybertechnology, or a feature of that technology, that is controversial from a moral perspective.

1a.   *Disclose any hidden or opaque features.*

By "hidden or opaque features" we mean aspects of the technology that have moral implications, where those affected by these aspects are either: not aware of them, or aware of them but not aware they have moral implications.

    1b.   *Assess any descriptive components of the ethical issue via the sociological implications it has for relevant social institutions and socio-demographic groups.*

    1c.   *Analyse the normative elements of the issue: determining if there are any specific guidelines (social policies, professional ethical codes) that can help resolve the issue.*

Step 2: Analyse the ethical issue by clarifying concepts and situating it in a context.

    2a.   *Establish what policies may apply in the specific context of the ethical issue.*

    2b.   *If there is a policy vacuum, clear up any conceptual muddles that may affect policy setting in this context.*

         Resolving a policy vacuum can be difficult if there is a *'conceptual muddle'* about what a key aspect of technology means. Note that such muddles occur because IT is 'logically malleable' — it can be used in many different ways for many different purposes.

Step 3: Deliberate on the ethical issue.

    3a.   *Apply one or more ethical theories to analyse the issue.*

         Ethical theories can focus either from an 'Act' perspective (looking at the moral acceptability of an act without reference to any existing rules) or a 'Rule' perspective (looking at the moral acceptability of an act if it is evaluated by an existing rule).

         We shall be using two popular ethical theories: Utilitarianism, which focuses on the consequences of an act; and Deontology, which focuses on the means by which an act is carried out.

    3b.   *Argue a position and justify it.*

5   Analysis and Results

The following applies Tavani's CCM approach to examine the moral implications of different types of work rejection on the AMT platform. Note that Step 3b, discussing the implications of the findings from the analysis, is subsumed into the Discussion section of the paper.

5.1   STEP 1(a): Disclose any hidden or opaque features.

We begin our analysis by determining if there are any hidden, or opaque, elements embedded in the AMT technology that could have moral implications in terms of the rejection of Worker submissions.

First, there is the API function that provides the Requester with the capability to automate different aspects of the HIT process. The Workers may, or may not, be aware of this capability, but regardless, they will not be able to discover if a particular Requester is using it, and for what reason. As noted in Situation (v) above, the Requester may seek to increase efficiency by automating the evaluation function, and increase effectiveness by systematically checking on Worker attentiveness and verifying accuracy of the outcomes. But these API functions are reported to be of varying quality (Silberman and Irani, 2016), and can lead to systematic, unfair rejections. Thus the API represents a potentially morally opaque, unknown feature from the Worker perspective.

Secondly, the tasks described in a HIT list may not be wholly visible to the Worker when they preview a particular HIT, as the AMT system limits how much can be seen to one screen page. This means Workers may end up accepting a HIT that involves more effort than they expected. Thus while the technology feature may be known to Workers (e.g. from past experience), key aspects of the task may remain hidden, creating arguably unnecessary uncertainty for the Worker seeking to avoid work rejection, and thus may have moral implications.

5.2   STEP 1(b): Assess any descriptive components of the ethical issue.

There are three stakeholders involved when submissions to HITS are rejected: The Requester, the Worker, and Amazon who run the AMT platform. Here we consider both the general implications of

rejection for each of the three parties, as well as the moral implications involved with each of the six reasons for rejection identified above.

5.2.1   Implications for Workers

Rejected work has clear negative impacts for Workers. First, they will not be paid for any of the work they have put into the submission. Second, a rejection is recorded against them in AMT's reputation system which, in turn, would negatively affect their likelihood of offers for better paying HITS, and their ability to attain Master status on AMT. Too many rejections with the same HITS may see them blocked by the Requester (as might occur in Type A situations). Too many Requester blocks will lead AMT to assume the Worker is a consistently bad performer, and the Worker's account may be suspended.

There may also be an element of time and effort (possibly accompanied by frustration — more on this later) spent communicating with Requesters to avoid rejection (especially in Type B situations) or establish why rejection occurred (especially in Type C situations). Note that Type C situations, if not resolved, may also lead to the Worker complaining to AMT.

5.2.2   Implications for Requesters

Requesters do not lose money through rejected work: they do not have to pay the Worker, and no corresponding commission is paid to AMT. However, repeated rejections across different Workers may defeat the purpose of the HIT in the first place, unless the Requester uses the submitted work anyway, as would occur in Type C(vi) situations, and possibly some Type B situations as well. Further, too many rejections across many Workers will adversely affect the Requester's reputation, which is monitored on Worker online forums and non-AMT reputation systems like Turkopticon.

There is also an element of time and effort responding to Worker queries (as noted above), though it seems Requesters' willingness to respond to queries will vary.

5.2.3   Implications for AMT

Every time a HIT submission is rejected, AMT does not get paid the commission from the Requester. However, beyond that, there are few costs incurred. The majority of the activities conducted across the AMT platform are automated, and thus incur little overhead. Monitoring Requester blocks and making decisions on account suspensions may require human involvement, but these are unlikely to be common occurrences. Responding to complaints from Workers may also require some human involvement, but as indicated later in this paper, non-automated responses do not appear to be common in practice either.

5.3   <u>Step 1(c)</u>: Analyse normative elements of the issue

There are two possible areas that could help address issues around rejection of, and non-payment for, work. The first involves the existence of any industry or professional ethics guidelines; but this is a relatively new industry involving digitally enabled work that does not appear to fall under the ambit of any existing guidelines for ethical practice.

The second possibility is employment law. The role of employment law in the 'gig economy' has been a contentious issue for over a decade (Stone, 2006; Friedman, 2014). In this case, the legal situation focuses on the employment status of Workers who perform microtasking jobs for others through a crowdsourcing platform. AMT makes it clear in the Participation Agreement that Workers are considered to be independent contractors, and not employees. This means that Workers are not eligible for the rights and benefits that would normally accrue to employees. In terms of the rejection of submitted work, this would have little impact on Type A situations, and arguably little effect on Type B situations unless the sums of money involved justified some form of arbitration. Type C situations, particularly C(vi), which effectively constitutes wage theft, would be subject to legal redress if Workers were classed as employees (Felstiner, 2011).

5.4   <u>Step 2(a)</u>: Establish what policies might apply in this context

AMT does provide both policy and guidelines in relation to the rejection of work, and these can be found on AMT's website.

5.4.1   The role of AMT's policies

AMT policy regarding work rejection is straightforward. The Participation Agreement states that "If a Requester is not reasonably satisfied with the Services, the Requester may reject the Services." This is to safeguard Requesters from having to pay for poor quality work, and consequently provides incentive for Workers to strive for the quality that Requesters seek. But no indication of what constitutes 'reasonable' in terms of Requester satisfaction is provided. This effectively enables any of Type A, B or C rejections to occur while still complying with the policy.

Requesters are required to provide an explanation as to why work has been rejected. However, this is enforced by providing Requesters with a text field that cannot be left blank. This means Requesters may, in fact, enter anything they wish. Consequently, while some Requesters provide adequate explanations as to why work is rejected, many don't (Silberman and Irani, 2016). Consequently, apart from Type A(ii) rejections, Workers may be left no wiser as to why their work has been rejected. If this occurs while a HIT is still active, it means Workers will not have an opportunity to change the way they complete the work to improve the quality.

AMT makes it clear in the Participation Agreement that they have "...the right, but not the obligation, to monitor any activity, content and Materials associated with the Site. [AMT] may investigate any reported violation of its Policies or complaints and take any action that it deems appropriate." However, what this means in practice is very unclear, and there appears to be little evidence that action is taken by AMT where complaints, particularly from Workers, do occur (Silberman and Irani, 2016).

5.4.2   The role of AMT's guidelines

Relevant AMT guidelines are provided for both Requesters and Workers. Worker guidelines are provided in the document *Worker Web Site FAQ*[3]. AMT advises Workers to contact the Requester directly if a Worker has concerns about the rejection of their work by using the email based system AMT provides. No other advice in relation to work rejection is provided.

Requester guidelines are provided in the document *Requester Best Practices Guide*[4], which offers considerable advice on the design and evaluation of HITs. This includes detailed advice on making instructions clear and concise, and testing the HITs to help discover technical errors and confusing instructions, all of which would help reduce unfair rejection described in Type B(iii).

5.4.3   The role of AMT's communication system

As part of their policy enforcement and guidelines for good practice, AMT encourages Workers and Requesters to communicate with each other and with AMT, where problems occur. In practice, however, this has proved consistently problematic. Workers may seek to communicate with Requesters if they are struggling with the HIT (e.g. the instructions are not clear), or if, as noted earlier, they do not understand why their work has been rejected. Requesters are not required to respond, and often don't. As one Requester noted: "You cannot spend time exchanging email [with workers]. The time you spent looking at the email costs more than what you paid them. This has to function on autopilot as an algorithmic system…" (Silberman and Irani, 2016, p.12). AMT notes in the Partici-

---

[3] https://www.mturk.com/mturk/help?helpPage=worker

[4] https://requester.mturk.com/help

pation Agreement that if there are any disputes between Requester and Worker, AMT "...will not be involved in resolving any disputes between participants relating to or arising out of the Services or any transaction."

Similarly, any complaints from Workers relating to perceived unfair rejections of work or lack of explanations from Requesters for rejections are either ignored or engender 'form' responses. AMT simply remains disengaged.

5.5 Step 2(b): Clear up any conceptual muddles

The word 'rejection' in this context appears to be poorly conceptualised. We have noted there are six broad reasons that would explain why work is rejected by Requesters, and yet AMT has provided policy that addresses only type A rejections, by allowing Requesters to reject work as they see fit. There are guidelines, but no policy, to address Type B rejections. AMT provides nothing to recognise or address Type C rejections, including malicious C(vi) rejections.

5.6 Step 3(a): Apply ethical theories to analyse the issue

There are two theories outlined by Tavani that can be used to evaluate the moral questions arising from AMT's rejection system and the way it is used in practice. Each is considered in turn.

5.6.1 Evaluating the rejection system from a Utilitarian perspective

As noted earlier, there is no policy available to clearly distinguish between appropriate and inappropriate rejection of work. Therefore, applying Tavani's (2013) formulation of act utilitarianism to work rejection on the AMT platform leads to the following statement:

> The act of rejecting submitted work is morally permissible if the consequences produced by this act result in the greatest good for the greatest number of persons affected by this act.

Before assessing this statement, we need to consider what is meant by "greatest good" and "greatest number". As Tavani (2013) notes, in this case the concept of 'good' is based around the idea of 'social utility', or social usefulness, "rather than via abstract criteria such as individual rights or social justice" (p.54).

In terms of 'number', there are two ways to consider rejection. If we consider the act of rejecting an individual Worker's submitted work, then we note AMT has missed a relatively small piece of revenue, the Requester has lost only a small piece of time, whereas a Worker may have lost an important segment of income. A Type A rejection may lose the Worker money, but had the Requester accepted that work, they would have had to pay for inferior work that they possibly couldn't use, through no fault of their own. Therefore, a Type A rejection is morally permissible. Similarly, a Type C rejection, which should have been an acceptance, is not morally permissible. Unsurprisingly, Type B rejections may, or may not, be morally permissible, depending on both the level of the Requester's negligence in setting the task, their willingness to improve it when contacted, and the level of reasonableness shown by the complaining Worker.

However, the Utilitarian view becomes more interesting if we consider all the rejections that occur over the history of a particular HIT. If a HIT generates only a small number of rejections, then we could consider the argument over individual rejections above. However, where larger numbers of Workers have their work rejected for a particular HIT, we note that AMT would still lose only a relatively small amount of revenue, the Requester would have lost time, but not money, while a significant number of Workers would lose potentially important income. The arguments for this case are similar to those above for the single rejection case, but with two key differences. The potentially greater negative impact on a Worker becomes magnified when looking at a large number of Workers. Second, a large number of rejections for a HIT suggests a systemic problem with either the HIT itself (as might happen with Type B rejections) or the way the work is evaluated (as might happen with Type C rejections). In these cases, it can be argued that mass rejections are morally impermissible.

5.6.2   Evaluating the rejection system from a Deontological perspective

Applying Tavani's (2013) formulation of act deontology to work rejection on the AMT platform leads to the following statement:

> The act of rejecting submitted work is morally permissible if the act adheres to *prima facie* (self-evident) duties.

There are two principles that should apply here: universality and impartiality. In other words, these duties, or obligations, should apply to everyone, and should not be seen to favour any particular individual or group over others. Examples of *prima facie* duties include: honesty, benevolence, justice, etc.

In the case of work rejection, three key duties/obligations can be identified: honesty (of intent), fairness (of action), and responsibility (for governance and management). We consider each of the three stakeholders in turn.

The onus of honesty and fairness in terms of work rejection falls onto the Requester and the Workers. Type A rejections may be seen to be fair, while Type C rejections unfair. Type A(ii) and Type C(vi) rejections amplify these assessments due to the levels of dishonesty (by Workers and Requesters respectively) on display.

AMT does more than just provide an opportunity for Requesters and Workers to meet. Using the AMT platform set up for microtask transactions, they are able to control who participates, what transactions are permitted, how the transactions will be conducted, and how payment will be organised. This means AMT have considerable responsibility to both govern the microtasking process and manage its operations. Governance arrangements involve the creation of policy, allocation of decision-making rights, and the creation of mechanisms to facilitate the application of policy and decision-making (Johnstone et al., 2006). There is little information about AMT's governance arrangements beyond what is published on their website. Policy appears largely intended to protect AMT from any adverse legal action. Amazon appoints a vice president responsible for AMT, but little is known about the decision-making structure within AMT. Visible mechanisms are essentially software-based, operated through the platform. In terms of work rejection, AMT policy makes it clear they place the power of rejection with the Requester, and any disputes that ensue fall outside the responsibility of AMT. AMT offers limited mechanisms to help with problems that may occur. For example, if a Worker has a problem with their work being rejected, there are communication channels available to complain to the Requester and/or AMT. However, in practice, there is no requirement for either of these parties to respond appropriately, and evidence suggests they frequently don't (Silberman and Irani, 2016). AMT does not offer a mechanism to deal with Requesters who reject work on a larger scale. However, if a Worker has work rejected too often, there is a reputation mechanism that will ensure they are less likely to get future, better-paid work, and in extreme cases can be blocked from participation, or even have their account closed. Thus, AMT has governance arrangements in place that can deal with more extreme cases (in terms of frequency) of Type A rejections, which may be fair, but these arrangements can be seen as blunt instruments given they may apply equally to Type B and Type C rejections as well, which is demonstrably unfair.

6   Discussion

On the face of it, it would seem there is a balanced perspective on the nature of work rejection on the AMT platform. Type A rejections seem to be fair, Type C rejections seem to be unfair, and Type B rejections will, on a case-by-case basis, fall somewhere within the fair-unfair continuum. The question is to what degree the rejection system is organised to deliver a balanced outcome for the parties involved.

Most countries have workplace laws that have developed over decades of negotiation and have sometimes emerged from conflict between various stakeholder groups. These laws aim to maintain some kind of ethical balance between the interests of these groups. New forms of work, such as the type of digital work that takes place on Internet based platforms like AMT have not evolved in this context.

While some may argue that workers accept the terms when they accept tasks, others would suggest that an ethical balance of risks and rewards for various stakeholder groups should be maintained in the same way as we expect it to be in other workplace contexts.

Our ethical analysis suggests the rejection system does *not* reflect this balance, and is, in fact, skewed heavily in favour of the interests of AMT, and to a lesser degree the Requester, as opposed to the interests of the Workers. A focus on the consequences of rejection (the Utilitarian view) has seen disproportionate effects on the stakeholders. The system is designed to punish Workers who have their work rejected, first by ensuring they do not get paid for the work they have done, and second by lowering their official work reputation and making it harder to achieve Master status, regardless of whether rejection was appropriate or not. While Workers come from a range of backgrounds, with varying reasons for engaging with AMT, a recent major study (based on self-reporting) found that 25% of Workers regard AMT as their primary source of income (Hitlin, 2016). The same study found that 91% of AMT Workers earned less than $8 (US) per hour, which would suggest many Workers are precariously dependent on AMT for financial survival. This means many Type B, and all Type C, rejections can unfairly have serious implications for the Workers. The widespread effects are amplified if unfair rejection occurs at a systemic level, which may occur with some Type B rejections, and will occur with Type C rejections. Finally, while dishonest Workers (Type A(ii) rejections) are punished, it would seem dishonest Requesters (Type C(vi) rejections) are not.

A focus on the obligations of those who reject work, and those who govern the rejection process (the Deontological view) supports the view that *prima facie* duties are not fully engaged. It could be argued that most Requesters strive to be honest and fair, but through negligence or ignorance may still treat submitted work unfairly. There is also little incentive for Requesters to explain to Workers why work has been rejected.

Of greater concern is the role that AMT plays in terms of governing the process. For example, guidelines provided for Requesters to help them design tasks that are clear about what has to be done and what constitutes acceptance/rejection criteria may not be enough to avoid unfair Type C(v) rejections. Some issues affecting rejection may even be hidden from Workers, such as the limitations on how much information about a task may be visible to Workers, or the use of API software to automate the work evaluation and accept/reject process. Punishment systems, as noted earlier, are skewed against Workers. Communication systems enabling Workers to express concerns to AMT are available, but are largely ineffective as there is no incentive for AMT to respond. Policies are designed to primarily protect AMT's interests, and provide Requesters with the last word on work rejection.

While the ethical analysis presented in this research focuses on only component of the microtasking process employed by AMT, we believe it is indicative of wider issues and makes several contributions.

First, this study contributes to an area of information systems that has received only modest academic attention. Deng et al. (2016) represent one of the few examples in IS literature directly addressing the potential for exploitation of Workers using AMT. They provide recommendations for all three stakeholders: Workers need to organise to create a stronger, collective voice; Requesters need to be more aware of the ethical issues facing Workers, and correspondingly improve HIT design and communication; and AMT needs to ensure Workers are paid better, protected from scams, and be more fair and transparent with their processes and policy enforcement. These are worthy suggestions, and many of them have been proposed before, but they are essentially saying AMT, along with negligent, ignorant or miscreant Requesters, need to behave better. Given AMT has not, so far, responded to any past, similar criticisms, it seems unlikely Deng et al.'s (2016) paper will lead to any significant change.

Second, this has implications for broader considerations within critical ethics. Stahl (2008) notes that emancipation is a frequent topic in critical ethics in IS research, where "capitalist work structures not only enslave and alienate labourers, but they also systematically take away their ability to develop and prosper" (p.141). A rigorous ethical analysis of just one contentious component of AMT's operation provided strong support for the view that AMT has little appetite for taking responsibility for any exploitation of vulnerable workers that may occur, and is accountable to no one. It is unclear what path

digital workers using these platforms can take to "develop and prosper" in the future. There has been increasing activity looking at the rights of crowdsourcing workers from the European Union (Spencer and Joyce, 2016), but challenges around jurisdiction and enforcement remain.

On a different note, this discussion has a high degree of relevance close to home. AMT, and similar platforms are widely used for data collection to support research. Recent studies suggest academics are now among the largest group of Requesters on AMT (Deng et al., 2016; Hitlin, 2016). However, there is no evidence to suggest academics are more ethical in their relationship with Workers than other Requesters. Recently, there has been strong interest registered in the IS literature promoting the use of AMT for research (Steelman et al., 2014; Lowry et al., 2016; Jia et al., 2017), but this has not been accompanied by any guidance on doing so ethically. We would argue that this is the right time for the Association of Information Systems to encourage debate on this issue, perhaps with the aim of including provisions within their Code of Research Conduct[5]

7. Conclusion

Our ethical analysis suggests that the policies and processes governing rejection of work on AMT do not reflect an ethical balance or risks and rewards. Considered from a utilitarian perspective, the imbalance of risks can be seen to be skewed in favour of the requester and the AMT platform. Considered from a deontological perspective, *prima facie* duties such as honesty, benevolence, and justice can be subverted by a lack of accountability on the part of AMT for resolving disputes. While we do not, at this stage, propose specific changes to AMTs policies and processes that will address the ethical imbalance, the first steps towards improvement are awareness that a problem exists and clear diagnosis of the problem. Our analysis supports a more informed dialogue of the ethical issues involved in rejecting crowdwork (using AMT as an example). This, in turn, can be used to design more balanced, equitable, and ethical policies and processes.

References

Boudreau, K. and K. Lakhani (2013). "Using the crowd as an innovation partner." *Harvard Business Review* 91(4), 60-69.

Deng, X., K. Joshi and R. Galliers (2016). "The duality of empowerment and marginalization in microtask crowdsourcing: giving voice to the less powerful through value sensitive design." *MIS Quarterly* 40(2), 279-302.

Ettlinger, N. (2016). "The governance of crowdsourcing: rationalities of the new exploitation." *Environment and Planning A* 48(11), 2162-2180.

Felstiner, A. (2011). "Working the crowd: employment and labor law in the crowdsourcing industry." *Berkley Journal of Employment & Labor Law* 32(1), 143-202.

Friedman, G. (2014). "Workers without employers: shadow corporations and the rise of the gig economy." *Review of Keynesian Economics* 2(2), 171-188.

Harris, M. (2014). *Amazon's Mechanical Turk workers protest: 'I am a human being, not an algorithm'*. The Guardian. URL: https://www.theguardian.com/technology/2014/dec/03/amazon-mechanical-turk-workers-protest-jeff-bezos (visited on 11 Nov 2017).

Hitlin, P. (2016). *Research in the crowdsourcing age, a case study*. Pew Research Centre.

---

[5] http://c.ymcdn.com/sites/aisnet.org/resource/resmgr/Admin_Bulletin/AIS_Code_of_Research_Conduct.pdf

Jia, R., Z. Steelman and B. Reich (2017). "Using mechanical turk data in IS research: risks, rewards, and recommendations." *Communications of the AIS* 41(Article 14), 301-318.

Johnstone, D., S. Huff and B. Hope (2006). "IT projects: conflict, governance, and systems thinking." In: *Hawaii International Conference on Systems Sciences (HICSS)*. Hawaii.

Kavanaugh, S. (2017). *Virtual sweatshops paint a bleak picture of the future of work*. Vocativ. URL: http://www.vocativ.com/410794/are-virtual-sweatshops-the-future-of-work/index.html (visited on 9 Nov 2017).

Lowry, P., J. D'Arcy, B. Hammer and G. Moody (2016). "'Cargo cult' science in traditional organization and information systems survey research: a case for using nontraditional methods of data collection, including Mechanical Turk and online panels." *Journal of Strategic Information Systems* 25, 232-240.

McInnis, B., D. Cosley, C. Nam and G. Leshed (2016). "Taking a HIT: designing around rejection, mistrust, risk, and workers' experiences in Amazon Mechanical Turk." In: *ACM CHI Conference*. ACM. San Jose, USA. 2271-2282.

Pittman, M. and K. Sheehan (2016). "Amazon's Mechanical Turk a Digital Sweatshop? Transparency and accountability in Crowdsourced Online Research." *Journal of Media Ethics* 31(4), 260-262.

Silberman, M. and L. Irani (2016). "Operating an employer reputation system: lessons from Turkopticon, 2008-2015." *Comparative Labor Law & Policy Journal* 37(3).

Spencer, N. and S. Joyce (2016). *Crowdwork in Europe.* University of Hertfordshire.

Stahl, B. (2008). "The ethical nature of critical research in information systems." *Information Systems Journal* 18, 137-163.

Steelman, Z., B. Hammer and M. Limayem (2014). "Data collection in the digital age: innovative alternatives to student samples." *MIS Quarterly* 38(2), 355-378.

Stone, K. (2006). "Legal protections for atypical employment law for workers without workplaces and employees without employers." *Berkley Journal of Employment & Labor Law* 27(2), 251-286.

Tavani, H. (2013). *Ethics and Technology: Controversies, Questions, and Strategies for Ethical Computing*. USA: Wiley.

Williamson, V. (2016). "On the ethics of crowdsourced research." *PS: Political Science & Politics* 49(1), 77-81.