**Title**

Algorithms of Suspicion



**Author**

Irani, Lilly

**Publication Date**

2026-06-24



Peer reviewed

# Algorithms of Suspicion

Lilly Irani

## Introduction

Platform workers for whom jobs might be a lifeline could wake up one day to find themselves cut off from work without notice. The reasons might be a software glitch, a change of address, or myriad others as this chapter's analysis will reveal. For those workers brave or stubborn enough to contest it, they may be met by silence from platform technical support. They may be told that there is nothing the platform can do. Or, if they are lucky, the platform may investigate their case and determine their account was suspended in error. In those cases of repair, companies have not redressed their error by restoring earnings workers lost shut out of their accounts. Turkopticon, one platform worker advocacy project, assists many workers in this situation. The consequences are lost access to livelihood, already earned wages trapped in platform accounts, and a deep sense of frustration and unfairness. Workers report such experiences of automated firing, wage theft, and platform silence on myriad gig work platforms, including Uber, Lyft, Doordash, and, the focus of this chapter, Amazon Mechanical Turk (AMT).

This chapter builds on scholarship on automated shop floor management – a problem that includes but is also broader than platforms. Jeremias Adams-Prassl has examined how the introduction of algorithms into employer functions of hiring, monitoring, and firing concentrate managerial control over work while also making responsibility and accountability difficult to track.[i] Nantina Vgontzas explains how algorithms in Amazon's warehouse operations are central to controlling workers. They do this through surveillance and control of rates that produced gendered and racialized hierarchies in the workplace. The algorithms, engineered by Amazon's own tech workers, also direct packing and shipping to work around sites of work slowdowns or strikes, making it more difficult for workers to engage in effective, concerted collective action.[ii] More broadly, a wide ranging scholarship on gig work has debated platforms' forms of control through reputation systems, gamification, and debt, as well as platforms' strategies of legal and especially employment (mis)classification.[iii]

This chapter contributes to these debates by drawing attention to a form of algorithmic control and wage theft through account suspensions on one such platform, Amazon Mechanical Turk. Policy gives platforms wide berth to suspend under laws that exempt "fraud detection" activities from oversight and transparency frameworks. In many jurisdictions, workers are also (mis)classified as independent contractors, lacking rights of recourse in cases of firing. Finally, Amazon's algorithms – marketed to other companies as Amazon Fraud Detector – are also of consequence not only to its platform workers, but to small businesses and consumers on its wider website and to internet users at large. As goes Amazon, so goes the world?

This chapter draws as data the experiences and insights of AMT workers faced with the problem of account suspensions, analyses of how AMT's interfaces and policies structure and automate employer-worker interactions, and analyses of patents that give clues as to Amazon's black box fraud detection algorithms unleashed on workers at scale. I show legitimate reasons why workers might get caught up by Amazon's algorithms that catch workers whose behavior appears suspicious to an algorithm with no social or political knowledge of their lives. The chapter concludes by arguing that existing digital rights frameworks must be revised to give workers rights and protections against platforms' algorithmic forms of management. I offer preliminary suggestions to this end.

This is a work of engaged scholarship, accountable to workers on the frontlines experiencing these workplace issues at the intersection of policies, practices, and algorithms. The research is particularly indebted to the work of AMT workers organizing as Turkopticon, a project that advocates for better working conditions for workers. An important part of their work is assisting AMT workers whose contesting account suspensions. Workers' problems and insights motivate and inform this chapter's inquiry. Workers' comments on this paper and conversations have informed its development. This work is informed by methods such as Community-Based Participatory Research and Design Justice.[iv] From those methods, I adopt accountability to communities and a respect for communities' expertise in many circumstances that affect their lives.

Policy, practices, and design are mutually constitutive in complex systems design.[v] Workers encountering these problems propose solutions that might require shifts in one or more of these three aspects. Solutions to the problem of account suspensions likely requires changes in intermediary (e.g. platform, company, employer) practices, including algorithm design processes, how employers (e.g. platforms) employ algorithmic judgements, and how they prevent or repair unwarranted harm to workers from those judgements. Solutions may also require changes to the algorithms themselves, the data sets on which the algorithms are trained, or the ways in which outputs are applied. Finally, policy environments set the conditions in which companies configure these algorithmic processes, create risks for workers, and distribute the costs of those risks. This chapter analyzes AMT's fraud practices as co-constituted by organizational practices, algorithms, and a permissive legal environment that allows algorithmic managers wide berth to opaquely regulate workers without worker redress.

**Amazon Mechanical Turk: Classifying cultural data for AI**

Amazon Mechanical Turk (AMT) is a data work marketplace that powers machine learning (ML) and artificial intelligence (AI) pipelines across the tech industry.[vi] Amazon launched AMT in 2006 to allow programmers to issue data processing calls directly from their computer code. Rather than other code answering the call, thousands of workers wait at their computers, ready to perform cognitive piecework on demand. The system, and systems like it, are critical to the production of ML and AI.

AMT workers choose among tasks like transcription, content moderation, and image classification, getting paid per piece of data processed. They categorize and classify data so that machine learning algorithms can learn to approximate Turkers' judgements. Sometimes the workers help train AI. Other times, they make up for AI's shortcomings. A company that scans customer receipts into budget software, for example, would send the faded or irregular receipts that confused its algorithms to AMT workers to transcribe and enter into the software. AMT workers may also participate in social science experiments and data gathering as paid participants. AMT workers, in short, perform those tasks of cultural judgement they learn to perform through a lifetime of cognitive and cultural development.

Amazon first patented the system in 2007 to help organize products from different sellers on Amazon's marketplace.[vii] Engineers faced a problem: they had data, including pictures and text, about products from different merchants to display on their websites. A stream of products uploaded by merchants, of, say, pens may not have a standardized identification number. For example, are products titled Lamy Safari, Lamy Safari Pen, Lamy Safari Fountain Pen, and Lamy Fountain Pen Cartridge the same product? Computers cannot tell without human interaction. A person familiar with fountain pens might easily say that the first three are names for the same thing but the last is a different object, even though each phrase differs from the next by only one word. So, Amazon's engineers had a problem. How could they automate the process of creating just one page for a product and linking to different merchants that offer the product but describe it in varying ways?

This kind of problem is one of organizing cultural data – data that contains cultural categories produced and changed by people as they interact through language and symbols.[viii] Artificial intelligence, machine learning, platform content moderation are all examples of attempts to automate at least part of the process of organizing, filtering, and responding to cultural data. The engineer's goal is to extract correct categorizations from workers to train machine learning or substitute for it when it fails. Just as a lot of work goes into quantifying and scaling the messy world so it can be represented statistically, a lot of work also goes into both producing correctly classified data sets to produce these models, as well as correcting the answers of models when they fail. Correct here is employers who enforce dominant culture definitions and conventions through their task design and criteria.[ix]

On Amazon Mechanical Turk, this work is completed by over 120,000 registered workers, through only 30,000 were active in a given month according to one 2020 study.[x] A 2016 International Labour Organization (ILO) study found that 45.2% of US and 90.7% of Indian AMT workers had college or post-graduate degrees, belying assumptions that workers on the platform lacked sufficient education to gain higher paying work. The ILO study found that almost 85% of workers were based in the US, 15% were from India, and a small percentage were from other countries. Workers, the ILO found, came to the platform for many reasons, including to complement pay from other jobs, to combine work with child or elder care, and to earn a living with a disability.[xi]

**Recruiting 1,000 workers in a day**

Employers outsourcing data processing work to the AMT platform create batches of tasks, called "human intelligence tasks" or "HITs" in Amazon parlance. HITs are web-based forms, hosted on Amazon's platform, that specify an information task and allow workers to input a response. Tasks include structuring unstructured data (e.g. entering a given webpage into an employer's structured form fields), transcribing snippets of audio, and labeling an image (e.g. as pornography, or violating given Terms of Service). Employers define the structure of the data workers must input, create instructions, specify the pool of information that must be processed, and set a price.

The employer then defines criteria that candidate workers must meet to work on the task. These criteria include the worker's "approval rating" (the percentage of tasks the worker has performed that employers have approved and, by consequence, paid for), the worker's self-reported country, and whether the worker has completed certain skill-specific qualification exams offered on the platform. This filter approach to choosing workers, as compared to more individualized evaluation and selection, allows employers to request work from thousands of temporary workers in a matter of hours.

One consequence of seeing people at scale has been that engineers imagine a world of "customers" and "bad actors." In their public writings to one another, engineers worry about two kinds of "bad actors": workers using scripts to generate random answers and hoping to get away with it, as well as workers not competent in the cultural categories needed by requesters.[xii] Workers might also produce "errors" because of interface design, exhaustion, undesirable interpretations, or simply by accident. AMT's restrictive interface prevents employers from knowing the difference between good faith errors, good faith cultural mistranslation, and bad faith "spam" work in many cases.

The language of "bad actors" reflects a wider culture of security beyond Amazon Mechanical Turk. It is common for large computer system operators, from corporations to the military, to operate in an imaginative geography of "good" and "bad actors." In such a world, engineers often see themselves in warfare with "bad actors" trying to "game the system." The more "bad actors" know about how the system operates, the more they can game it – or so the "security by obscurity" story goes. Conversely, engineers charged with protecting systems use the data available to them to identify and block "bad actors" from accessing the system. In a world of platforms made into a cybersecurity battleground, collateral damage is common.

In 2018, the "bot scare" erupted in the world of Amazon Mechanical Turk workers and employers, illustrating how employers imagine mundane workplace challenges as threats by "bad actors" or quasi-criminals. The scare began when a graduate student posted to a Facebook group of other Amazon Mechanical Turk employers observing that he was seeing more surveys with nonsense answers and respondents with identical locations. By the end of the week, researchers all over the internet were reporting similar concerns, fearing that their published research papers had analyzed spurious data produced by automated scripts – automated scripts animated by workers defrauding well-meaning researchers and breaking Amazon's terms and conditions along the way. Crowdsourcing research firm CloudResearch produced a study involving a range of tasks meant to discern whether answers were produced through automation or human input.[xiii] They reported that odd results were not bots at all, but rather workers who performed well on certain questions and poorly on others. These workers, often in south Asia, performed worse on certain kinds of tasks not designed with their capabilities in mind, but actually performed better than US workers on tasks geared towards their cultural knowledge. Requesters' panic about fraudulent workers in fact masked their lack of awareness of who their workers were, what their skills were, and what tasks they might be suited to. Patents and reported practices reveal a range of ways engineers attempt to discern "bad actors" from workers they ought to trust. They do this both at the scale of assessing work products, or HIT results, and also at the scale of assessing workers.

*Assessing work products*

Engineer-employers need to authenticate answers from untrusted workers to common data classification tasks that provide fuel (training sets, algorithmic evaluations) for machine learning. The engineer can extract answers from workers, but how can they know which answers are correct by their standards? Remember they are working with large volumes of data work output so they can't review answers individually. And they are dealing with socially negotiated categories and meanings. Is this porn or not porn is only the most extreme example. I have not yet seen an engineer grapple with anthropological questions of whether multiple answers can be "correct" in different contexts.

Engineers have several ways of dealing with this uncertainty. Sometimes employers will have a trusted "gold standard" data set of "ground truth" answers — correctly classified cultural symbols. These data sets might be produced and validated by AMT (or similar) workers, or acquired from trusted projects, labs, or other public datasets (Stanford, Iriondo, and Shukla 2020). They can use this "gold standard" to test workers, slipping in tasks with known correct answers into a stream of tasks with unknown correct answers. Another strategy, detailed in a patent but also commonly discussed by engineers, can be summarized as "the most plural judgements".[xiv] This can mean simply assigning several workers the same task and using majority vote to decide on the "true" answer.

With both the "gold standard" test and the "plurality agreement" test, engineers face a tradeoff: the more "gold standard" test tasks or duplicated tasks, the greater the certainty that the workers giving the correct answers are actually trustworthy rather than lucky. But the more test tasks employers pay for, the greater the "human capital investment".[xv] In patents and trade show presentations alike, engineers treat the confidence level needed in a worker a kind of knob that can be dialed up or down adding more or less information about workers.[xvi] This additional information can come by testing workers or gathering other information about them, such as through surveillance. The latter I will return to.

*Assessing workers*

Assessing human *judgements* can be a lot of extra work for engineers who have to put out multiple instances of a task, compare the answers, and figure out which workers to distrust or even blacklist. To reduce this work, Amazon has developed a program it calls "Masters." Amazon promotes "Masters" workers as trusted producers of high-quality work and charges a premium for them. Masters workers do not earn a higher wage per task, but they get exclusive access to a large stream of work. Amazon makes "Masters" workers a default for employers setting up tasks on the platform.

Workers find the Masters designation puzzling and frustrating. There is no application process or clarity about the criteria to achieve this. You wake up one day and find out you've been chosen as a Master. On forums, workers will talk to each other to compare their statistics and try to decode why some get the designation and some do not. Excellent performance by the simple statistics workers get on the platform – the percentage of approved work tasks they've done, the number of tasks they've done – do not seem to predict getting the designation. What is interesting here, though, is that Amazon has developed a meritocracy among workers and its mechanisms are intentionally opaque, likely to prevent "gaming" by those imagined "bad actors." (This is a world in which studying for the test is considered cheating.)

A 2011 Amazon patent holds some clues. The patent suggests that Amazon may calculate a "judge error rate" or confidence rating it assigns to workers behind the scenes (Figure 1).[xvii] Amazon has an advantage over engineer-employers in that it can maintain a record of the workers' error rates across their whole work history on the platform, across all employers. The patent suggests that Amazon, like engineers I've described, might use the kinds of "gold standard" and plurality techniques I described to evaluate "judges." The patent suggests a wide range of qualities by which a worker may be evaluated, including accuracy, error rate, and speed of task performance.[xviii]

While these may seem straightforward attributes of competence, the algorithm is agnostic to conditions that explain the differences, such as disabilities or multi-tasked care work in the home. Turk workers work in varied conditions, in their homes, at the library, while at school. As independent contractors, this kind of flexibility is touted as precisely part of the appeal of gig work. But as Amazon judges the "judges", these variations in bodies, environment, and work process manifest to the algorithm as differences in speed or even interpretive variations that manifest as "error rate." Together, these can mean the difference between being algorithmically judged a Master or perhaps a fraudster.

*An exception: finding a trusted pool*

Once they have identified workers they trust, some engineers will develop custom lists of worker IDs they will clear for their tasks and treat this group as a trusted pool (Gray & Suri 2019). Some might create a Slack channel to communicate with these workers or participate in worker-run forums to engage these workers through informal channels. These communication forums allow workers and employers to engage in more dialog about the labor process and how to adjust it. These more stable relationships might seem to indicate a tendency, as described by Ronald Coase, for firms to form as a way of decreasing transaction costs of operating on the open market (1937). These trusted pool formations reveal a tension, as argued by anthropologists Ilana Gershon and Melissa Cefkin, between neoliberal autonomy and transactional efficiencies.[xix] The emergent organizational forms reveal one way that workers and employers cobble together platforms and software, usually leaving platform operators like Amazon untouched, to negotiate these tensions in organizing a less wasteful and fairer labor process.

Such employers, from a workers' perspective, are the exception. I show this exceptional form of collaboration to underscore how work organization could be otherwise from the approach to filtering and testing workers *en masse*. Why is it exceptional? Elsewhere, I've argued that part of the appeal of Amazon Mechanical Turk for engineers is that their managerial relationships and responsibilities are displaced into the technological platform (Irani 2015). Engineers describe feeling like the platform is "magic" because it allows them to focus on the coding work they are attached to, rather than the enactment of managerial authority that violates their image of Silicon Valley work cultures. Most employers want to keep their workers across the programming interface, out of sight to be filtered, transacted, and extracted.

**Enforcing one login, one body**

These techniques of assessing workers all depend on a one-to-one relationship between an Amazon worker login and the working body behind the screen. A login identifies the data object that holds the work history. When Amazon algorithms detect a possible violation of their assumption, they label accounts as at risk of being fraudulent. Workers only discover this when they get their accounts suspended.

One practice that Amazon penalizes is the sharing of a single worker login among several workers. This is a problem especially for Indian workers who often face barriers getting an AMT login. AMT does not accept all people who apply to be workers and, true to form for black boxed fraud governance, does not explain why they reject applications. It may be because Indian workers are less likely to be Amazon customers and thus lack a digital history, or it may be because Indian workers were more likely to be rejected by requesters for giving undesirable task responses and thus machine learning algorithms have learned to discriminate against workers from India in general. Some Indian workers unable to get their own account respond by sharing logins with others, working around the clock and dividing up the wages that accrue to the account. Amazon seeks to shut down this behavior that violates the integrity of one login, one body. Workers may suddenly find their accounts suspended, unable to access the previous earnings already held in their Amazon account.

Algorithms policing violations to one login, one body sometimes shut down practices that Amazon later determines, under pressure from workers, to be legitimate. In one case, a mother and son found their logins suspended because they were logged in from the same house, and thus the same wifi router, and, thus, had the same IP address. The accounts were only restored after Turkopticon worker organizers intervened to bring the case to an Amazon manager. Amazon explained their policy as a way of protecting the authenticity of social science research results. Amazon had interpreted the second login on the same IP address as a worker trying to get paid to do the same social science study twice, corrupting research results. While Amazon did admit its mistake when it reinstated the mother whose son turked from her home; for two weeks of lost earnings, they gifted her a $10 Amazon gift card far less than what she would have earned.[xx]

Enforcing one login, one body means platform operators scan data flows in suspicion that one login hosts many bodies, as well as suspicion that multiple logins can host one body. Both are a problem for political economies of knowledge extraction. As Amazon's algorithms automate fraud judgements, they may learn to simply exclude workers from Indian IP addresses, or IP addresses from Starbucks, or IP addresses from libraries in low-income neighborhoods without, as Burrell and Jonas argue, "understanding the systematic social and political conditions that produce differential behaviors online."[xxi]

**The World's Largest Fraud Detector?**

The techniques described above are of consequence not only to Amazon Mechanical Turk employers and workers, but to internet users at large. While Amazon's approaches to fraud remain opaque to outsiders (and likely even many insiders), the company has made its cloud-based Fraud Detector available as a platform service it sells to other organizations. This is significant for several reasons. First, it means Amazon disseminates its techniques to the internet at large. Second, Amazon becomes a central broker of data signals on users and can leverage data and risk assessments made for one client to inform assessments for another client. Amazon Fraud Prevention team lead Ryan Schmiedl described this bleed of risk data across sites as a "lift" clients get from Amazon: "we're taking patterns from repeated bad actors...so you're benefiting not just from what you know, but from what we know".[xxii]

Amazon's Fraud Detector might label as higher risk whose self-declared country and customer address are different, or whose IP address country and phone numbers are different. A user facing precarity and moving across state and country borders, as is common between San Diego and Tijuana, here becomes suspect. Similarly, someone couch surfing in a different location than their listed address might be flagged. The machine learning algorithm, according to Schmiedl, also flags practices that are simply outliers "unlike anything it has seen from legitimate customers."[xxiii] To be risky, then, is not to demonstrate threatening behaviors. It might simply be that one seems dissimilar to others represented in the dataset.

Fraud reduction can justify creative and expansive forms of surveillance. One patent (Figure 2) describes a data-thirsty technique that tracks user behavior habits, such as sequences of applications opened in the morning, combined with contextual markers such as location, microphone data, and relationship to surrounding objects. That's the "behavioral data collector" indicated in the patent figure. The system authenticates users when they are acting habitually, as discerned by the algorithm. When use patterns change, the system forces reauthentication.[xxiv] The patent claims an expansive range of surveillance modalities to protect the sanctity of digitized finance systems and firm financial flows. This patent also appears perfectly aligned with Amazon's fight against people consensually sharing a single internet login.

Penalties for those informally labeled as "fraudsters" include loss of work, accrued wages, and even login-restricted services across the web. As a login infrastructure provider, Amazon controls not only access to its sites but access to all kinds of other sites across the internet.

Fraud risk techniques do not always present users with clear consequences. It can, according to Amazon presentations, result in "friction." "Friction" can take many forms. It can be those CAPTCHAs that ask you to mark all the trucks, bridges, or stop signs in a set of images. It can come as requests for additional authenticating data (address, phone number, social security number). It can come as a temporary delay pending a customer service call. Friction allows Amazon to titrate caution against fraudsters against the desire to allow customers to pay them. When dealing with customers who purchase products, Amazon needs to balance potential costs of fraudulent activity against the costs of slowing or blocking revenue-generating customers. But when dealing with workers in Amazon Mechanical Turk, Amazon faces a surplus of interested workers and even turns applicants away.[xxv] Thus, it can aggressively suspend long time workers or even block newer, untried workers with little consequence to the company.[3]

## Governing algorithmic judgement

Scholarship on algorithmic judgement has identified several distinct kinds of problems, each of which imply distinct resolutions or mitigation strategies. Amazon Mechanical Turk workers facing reputational harm and account suspensions can face problematic data, problematic engineering assumptions, problematic algorithmic inferences, or the criminalization of practices that are inconvenient for company business models. While the first two kinds of problems have been discussed widely in scholarship, the latter two have attracted less attention among those concerned with workers' or consumers' rights.

### *Rights to see and repair data about platform users*

The simplest kind of problem faced by workers confronting account suspensions is the correction or updating of data. Workers, like financial service users, might be flagged when they log in from a location distinct from the address registered to their account. Unlike financial service users, however, workers rarely have the opportunity to set an updated location status on the platform to avoid getting flagged.

Legal scholar Frank Pasquale has called for policies that guarantee the right to review data records about oneself. The fraud algorithms that are the central concern of this chapter are a key element, I argue, of what Pasquale calls “the black box society” – a society in which opaque and automated surveillance and reputational judgements restrict one’s access to resources and activities.[xxvi] The right to review and correct data ensures one mechanism for contesting unfair judgements. Companies can be expected to resist such transparency, arguing that it would allow bad actors crucial tactical knowledge about the cybersecurity battlefield. However, cybersecurity researchers argue that a truly secure system is one that is robust against an adversary who understands how it is constructed.[xxvii] Transparency, conversely, improves security by expanding the community of people who can point out problems in the system and suggest mechanisms for repair.

Evidence from this chapter shows how the term ‘bad actors’ occludes unjust consequences of platform operator design, algorithms, and policies. For workers and users confronting algorithmic judgements, account suspensions exemplify an emergency – a moment when technical and social conditions create an unusual event. These events do not happen every day or even to most users. For this reason, data transparency ought not to place the responsibility for maintaining correct data on workers, adding an extra labor to the time they spend producing. Rather, it ought to be one means of mitigating harm in a larger process of repair when algorithms, like people, inevitably misjudge the social world and fail.[xxviii]

*Beyond ethics education: oversight and accountability to those impacted*

In some cases, patents reveal engineering assumptions that explicitly encode a model of knowledge or ability that reproduces hierarchies among difference. Recall, for example, the patent that assumed speed as one indicator of worker quality, ignoring differences in competing obligations, pacing preferences, or body mobilities. Such an assumption seems especially unjustifiable on a platform in which workers are paid piece wages, rather than per hour.[xxix] Such an assumption is also especially problematic given the propensity of computer work to produce repetitive stress injuries that might slow workers down.

Simply educating engineers in ethics or social science methods cannot address such gaps, though such education may make engineers more aware of the gap and the need to address it. Social justice approaches have focused on such assumptions and argued that those who design must do in ways accountable to those directly affected by the design.[xxx] Drawing on a wide range of scholarship in feminist studies, disability studies, and the knowledge of social movement actors, such an approach draws on the knowledge of differently situated actors about the emergent effects designs may have on their lives. Technology transparency and oversight policy approaches require companies to reveal details of algorithms and data practices that affect users; they also create mechanisms of public regulatory oversight over platform operators, giving users one lever for understanding, redress, and advocacy.

These transparency and oversight approaches must not exempt fraud algorithms from rights and oversight claims, but some currently do. For example, California Assembly Bill AB-1790 sets out requirements for how online marketplaces such as E-Bay or Amazon resolve user disputes with what the law calls "minimum fairness," including specifying in writing grounds for suspensions. Yet the law exempts the disclosure of *"information that would hinder any investigation or prevention of deceptive, fraudulent, or illegal activity."* This exemption gives broad berth for platform operators to deny transparency and oversight to users. Another example of policy that exempts fraud algorithms was the California Privacy Rights Act (CPRA) of 2018, passed by the legislature, which required businesses to respond to consumer requests to delete their personal information. This law also offered exemptions for consumer data kept to detect "fraudulent…activity."

*Shield platform users from systemic risks*

The problem of algorithms, Louise Amoore argues, is not only "the power to perceive, to see, to collect, or to survey a vast data landscape, but the power to perceive and distill something for action".[xxxi] When algorithms of suspicion target a user for outlier behavior or behavior that pattern matches prior examples of bad actors fed to the system, that person might find themselves snatched out of an ocean of users for reasons of correlation rather than observed malfeasance.

This algorithmic power to focus and make actionable based on correlation introduces a kind of systemic risk – a expansion of vulnerabilities as systems become more interconnected and complex.[xxxii] Algorithms draw on diverse sources of data through diverse modeling techniques being deployed and tested on people in their real lives, rather than in the lab. I argue that this generates an accumulation of risks in the platform management system, but the consequences of accumulated risk then is concentrated on those workers who find themselves suspended by mistake.

We need policies, norms, and practices that account for this risk and protect people from the harms such risks can generate. A simple example for Amazon Mechanical Turk workers might be estimating the lost wages from a mistaken account suspension based on past earnings and adding compensation duress and for time spent working towards resolution. Actual solutions should be developed in collaboration with those most directly impacted by these risks.

*Refuse simple and criminalizing claims of fraud*

Algorithms of suspicion are not only harmful because their inferences or data are problematic. They can also be harmful by scanning for and targeting behaviors that might not be morally wrong but are inconvenient, undesirable, or costly for those who profit through algorithms.

To illustrate how socially acceptable practices might be inconvenient for and thus targeted by companies, I offer an example from the history of Google. When I worked for Google in the early 2000s, it was common for ad agencies to have a single login shared among staff to bid on Google AdWords keywords and craft ads. When Google transitioned the AdWords system onto the master Google Accounts system, we forced each agency worker to create their own login. Google wanted to impose an identity between login and person in order to collect individualized data and create profiles for "personalized search" and personally targeted advertisements. (Prior to this, Google ads were contextual, based on what you searched for at that specific moment.)

Like Google, Amazon imposes an authenticated vision of an individual subject, penalized for interdependence with others and sharing of resources. The "wisdom of crowds" relies on the people as sensors of their environments, and requires — in fact, enforces — their independence.[xxxiii] It recasts collaboration – between mother and son, or workers in a cowork space – as collusion.[4]

Part of the work of this chapter has been to denaturalize fraud and its associated discourse of "bad actors." The meaning of fraud has shifted historically, from concerns about forged documents leading to improper property claims[xxxiv] to concerns about compromising investors when uncertain statements about the future cross over into lies.[xxxv] The language and techniques of fraud I have traced here are less concerned with whether a person has documentary authority or is truthful, but rather whether they can be trusted as a cultural worker and, in turn, as an authority on legitimate cultural meanings. A worker who offers a less "popular" answer or an answer the employer does not consider sensible can be labeled fraudulent, rather than simply culturally different. On platforms, as in domains of public services in the United States, accusations of fraud can serve the goals of accusing organizations to evade responsibilities as part of a broader social contract.[xxxvi]

## Conclusion

What does it mean that fraud regimes are now being imposed on everyday communicational and work transactions, suppressing or even criminalizing other systems of ethics and truth such as care, cooperation, or infrastructure sharing? Further, Amazon is generating a market for fraud detection as a service, spreading practices drawn from finance, internet security, and international relations with their anxieties about bad actors who must be kept in the dark. In generating a market, Amazon generates desire and extends its infrastructures into new arenas of social life. In describing fraud detection systems, engineers easily slip into deploying quasi-criminalizing language that re-narrate a fraud risk as an individual as malicious and untrustworthy. Such an orientation forecloses avenues for recognizing harm to workers and other users, recognizing badly designed platform workflows, or repairing problematic data and algorithmic inferences. Finally, algorithms of suspicion generate systemic risk unaddressed by ethics approaches that focus on transparency and accountability. Such approaches do not account for how algorithmic management generate new kinds of uncertainty and risk and then concentrate impacts on the shoulders of just a few. We need policies, norms, and practices that account for such risks and protect people from the harms such risks can generate and offer redress and compensation when an algorithm shuts them out.

**Acknowledgements**

I am indebted to helpful readings and discussions with scholars at AI Now, UCSD's Critical Data Studies group, participants in an SSRC workshop on authenticity organized by Wendy Chun, linguistic anthropologists at the Michicagoan Conference, graduate students at Cornell University, and the Platform Governance volume editors. I also thank Stuart Geiger, Ilana Gershon, and Nick Seaver, and the editors of this volume for reading and discussing drafts of this work.

**References**

Amoore, Louise. 2020. Cloud Ethics: Algorithms and the Attributes of Ourselves and Others. Durham, NC: Duke University Press.
AB-1790. https://leginfo.legislature.ca.gov/faces/billCompareClient.xhtml?bill_id=201920200AB1790
Collier, Stephen J., and Andrew Lakoff. "Introduction: Systemic Risk." *Limn* 1, no. 1 (2011). https://escholarship.org/uc/item/0j88z94p.
Costanza-Chock, Sasha. *Design Justice*. Information Policy. Cambridge: The MIT Press, 2020. https://doi.org/10.7551/mitpress/12255.001.0001.
Dubal, Veena. "The Time Politics of Home-Based Digital Piecework." *Center for Ethics Journal: Perspectives on Ethics*, July 4, 2020. https://doi.org/10.2139/ssrn.3649270.
Ekbia, H. R., and Bonnie A. Nardi. 2017. Heteromation, and Other Stories of Computing and Capitalism. Acting with Technology. Cambridge, MA: MIT Press.
Eubanks, Virginia. Automating inequality: How high-tech tools profile, police, and punish the poor. St. Martin's Press, 2018.
Felstiner, A. 2012. "The Weakness of Crowds." *Limn*.
Graeber, David. "Dead zones of the imagination: On violence, bureaucracy, and interpretive labor: The Malinowski Memorial Lecture, 2006." *HAU: journal of Ethnographic Theory* 2, no. 2 (2012): 105-128.
Gray, Mary L., and Siddharth Suri. Ghost work: How to stop Silicon Valley from building a new global underclass. Eamon Dolan Books, 2019.
Hamraie, Aimi. "Designing Collective Access: A Feminist Disability Theory of Universal Design." *Disability Studies Quarterly* 33, no. 4 (2013).

Hullender, Gregory N. 2014. Evaluation of task judging results. United States US8868471B1, filed September 21, 2011, and issued October 21, 2014. https://patents.google.com/patent/US8868471B1
Ipeirotis, Panos. 2017. "Why Was My Amazon Mechanical Turk Registration Denied?" January 17, 2017. https://www.behind-the-enemy-lines.com/2017/01/why-was-my-amazon-mechanical-turk.html.
Irani, Lilly. 2015. "The Cultural Work of Microwork." New Media & Society 17 (5): 720–739.
Irani, Lilly. 2019. Chasing Innovation: Making Entrepreneurial Citizens in Modern India. Princeton, NJ: Princeton University Press.
Jackson, Steven J., Tarleton Gillespie, and Sandy Payette. "The Policy Knot: Re-Integrating Policy, Practice and Design in Cscw Studies of Social Computing." In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 588–602. CSCW '14. New York, NY, USA: ACM, 2014. https://doi.org/10.1145/2531602.2531674.
Kauffman, K. 2020. Personal conversation. October 17.
Masterman, Michael Frederick. 2018. Authentication and fraud detection based on user behavior. United States US10108791B1, filed March 19, 2015, and issued October 23, 2018. https://patents.google.com/patent/US10108791B1
Moss, Aaron, and Leib Litman. "After the Bot Scare: Understanding What's Been Happening With Data Collection on MTurk and How to Stop It." CloudResearch, September 18, 2018. https://www.cloudresearch.com/resources/blog/after-the-bot-scare-understanding-whats-been-happening-with-data-collection-on-mturk-and-how-to-stop-it/.
Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press, 2015.
Prop 24. https://vig.cdn.sos.ca.gov/2020/general/pdf/topl-prop24.pdf
Roach, Levi. *Forgery and Memory at the End of the First Millennium*. Princeton University Press, 2021.
Schmiedl, Ryan, Nick Tostenrude, and Kara Suro. 2020. "AWS Re:Invent 2019: [REPEAT 1] Fraud: How to Detect and Prevent It Using ML (AIM322-R1)." August 31. https://www.youtube.com/watch?v=ZkPQtjETrBw.
Slupska, Julia, Jeanette Lowrie, Lilly Irani, and Deian Stefan. "How Secrecy Leads to Bad Public Technology." (2021). https://escholarship.org/content/qt89d7826n/qt89d7826n.pdf
Stanford, Stacy, Roberto Iriondo, and Pratik Shukla. 2020. "Best Public Datasets for Machine Learning and Data Science." Medium (blog). August 28, 2020. https://medium.com/towards-artificial-intelligence/best-datasets-for-machine-learning-data-science-computer-vision-nlp-ai-c9541058cf4f.
Sunder Rajan, Kaushik. 2006. *Biocapital: The Constitution of Post-Genomic Life*. Durham, NC: Duke University Press.
Bottomley, Emlyn. "Data and Algorithms in the Workplace: An Overview of Current Public Policy Strategies." Working Paper, Technology and Work Program. UC Berkeley Labor Center, November 17, 2020.
"Top 10 Principles For Workers' Data Privacy and Protection." Nyon, Switzerland: Uni Global Union, November 2, 2017. https://uniglobalunion.org/report/principles-for-workers-data-rights/.
Vgontzas, N. (2022) 'Toward Degrowth: Worker Power, Surveillance Abolition, and Climate Justice at Amazon', *New Global Studies*, 16(1), pp. 49–67. Available at: https://doi.org/10.1515/ngs-2022-0008.
Walsh, Nicholas, and Robert Zhu. 2020. AWS What's Next Ft. Amazon Fraud Detector. Accessed October 16, 2020. https://www.youtube.com/watch?v=FnFckqGa7XE.

Wallerstein, Nina, Bonnie Duran, John G. Oetzel, and Meredith Minkler. *Community-Based Participatory Research for Health: Advancing Social and Health Equity*. Newark, UNITED STATES: John Wiley & Sons, Incorporated, 2017. http://ebookcentral.proquest.com/lib/ucsd/detail.action?docID=5097167

---

[i] Adams-Prassl, Jeremias. "What If Your Boss Was an Algorithm? Economic Incentives, Legal Challenges, and the Rise of Artificial Intelligence at Work Automation, Artificial Intelligence, & Labor Law." *Comparative Labor Law & Policy Journal* 41, no. 1 (2021 2019): 123–46.

[ii] Vgontzas, Nantina. "Toward Degrowth: Worker Power, Surveillance Abolition, and Climate Justice at Amazon." *New Global Studies* 16, no. 1 (April 1, 2022): 49–67. https://doi.org/10.1515/ngs-2022-0008.

[iii] Rosenblat, Alex, and Luke Stark. "Algorithmic labor and information asymmetries: A case study of Uber's drivers." *International journal of communication* 10 (2016): 27; Dubal, Veena B. "Wage slave or entrepreneur?: Contesting the dualism of legal worker identities." *California Law Review* (2017): 65-123; Prassl, Jeremias. *Humans as a Service*. Oxford University Press, 2017; Doorn, Niels van, and Julie Yujie Chen. "Odds Stacked against Workers: Datafied Gamification on Chinese and American Food Delivery Platforms." *Socio-Economic Review* 19, no. 4 (October 1, 2021): 1345–67. https://doi.org/10.1093/ser/mwab028.

[iv] Wallerstein, Nina, Bonnie Duran, John G. Oetzel, and Meredith Minkler. *Community-Based Participatory Research for Health: Advancing Social and Health Equity*. Newark, UNITED STATES: John Wiley & Sons, Incorporated, 2017; Costanza-Chock, Sasha. *Design Justice*. Information Policy. Cambridge: The MIT Press, 2020.

[v] Jackson, Steven J., Tarleton Gillespie, and Sandy Payette. "The Policy Knot: Re-Integrating Policy, Practice and Design in Cscw Studies of Social Computing." In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 588–602. CSCW '14. New York, NY, USA: ACM, 2014. https://doi.org/10.1145/2531602.2531674.

[vi] Irani, Lilly. "The Cultural Work of Microwork." *New Media & Society* 17, no. 5 (2015): 720–39. https://doi.org/10.1177/1461444813511926; Irani, Lilly. "Justice for Data Janitors." In *Think in Public*, 23–40. New York Chichester, West Sussex: Columbia University Press, 2019. https://doi.org/10.7312/marc19008-003; Gray, Mary L., and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt Publishing Company, 2019.

[vii] Harinarayan V, Rajaraman A, Ranganathan A (2007) *Hybrid machine/human computing arrangement*. Patent no. 7,197,459. USA.

[viii] Winograd, Terry, and Fernando Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Reissue. Addison-Wesley Professional, 1987.

[ix] Miceli, Milagros, and Julian Posada. "The Data-Production Dispositif." arXiv, May 24, 2022. https://doi.org/10.48550/arXiv.2205.11963.

[x] https://www.cloudresearch.com/resources/blog/mechanical-turk-data-five-years-in-five-figures/

[xi] Berg, Janine. "Income Security in the On-Demand Economy: Findings and Policy Lessons from a Survey of Crowdworkers." *Comp. Lab. L. & Pol'y J.* 37 (2015): 543.

[xii] Moss, Aaron, and Leib Litman. "After the Bot Scare: Understanding What's Been Happening With Data Collection on MTurk and How to Stop It." CloudResearch, September 18, 2018. https://www.cloudresearch.com/resources/blog/after-the-bot-scare-understanding-whats-been-happening-with-data-collection-on-mturk-and-how-to-stop-it/.

[xiii] CloudResearch. "How Many Amazon Mechanical Turk Workers Are There in 2019?," June 25, 2019. https://www.cloudresearch.com/resources/blog/how-many-amazon-mturk-workers-are-there/.

[xiv] Irani, Lilly. "The Cultural Work of Microwork." *New Media & Society* 17, no. 5 (2015): 720–39. https://doi.org/10.1177/1461444813511926.

[xv] Schmiedl, Ryan, Nick Tostenrude, and Kara Suro. "AWS Re:Invent 2019: [REPEAT 1] Fraud: How to Detect and Prevent It Using ML (AIM322-R1)." August 31, 2020. https://www.youtube.com/watch?v=ZkPQtjETrBw.

[xvi] Ibid; Hullender, Gregory N. Evaluation of task judging results. United States US8868471B1, filed September 21, 2011, and issued October 21, 2014.

[xvii] Hullender, Evaluation of task judging results

[xviii] Ibid.

[xix] Gershon, Ilana, and Melissa Cefkin. "Click for Work: Rethinking Work through Online Work Distribution Platforms | Ephemera." *Ephemera: Theory & Politics in Organization* 20, no. 4 (n.d.): 103–29.

[xx] Kauffman, K. 2020. Personal conversation. October 17.

[xxi] Jonas, Anne, and Jenna Burrell. "Friction, Snake Oil, and Weird Countries: Cybersecurity Systems Could Deepen Global Inequality through Regional Blocking." *Big Data & Society* 6, no. 1 (January 1, 2019): 2053951719835238. https://doi.org/10.1177/2053951719835238.

[xxii] Schmiedl, Tostenrude, and Suro, "AWS Re:invent"

[xxiii] Ibid.

[xxiv] Masterman, Michael Frederick. Authentication and fraud detection based on user behavior. United States US10108791B1, filed March 19, 2015, and issued October 23, 2018.

[xxv] Ipeirotis, Panos. 2017. "Why Was My Amazon Mechanical Turk Registration Denied?" January 17, 2017. https://www.behind-the-enemy-lines.com/2017/01/why-was-my-amazon-mechanical-turk.html.

[xxvi] Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press, 2015. For policy calls to allow such redress for workers, see Bottomley, Emlyn. "Data and Algorithms in the Workplace: An Overview of Current Public Policy Strategies." Working Paper, Technology and Work Program. UC Berkeley Labor Center, November 17, 2020; see also "Top 10 Principles For Workers' Data Privacy and Protection." Nyon, Switzerland: Uni Global Union, November 2, 2017. https://uniglobalunion.org/report/principles-for-workers-data-rights/.

[xxvii] Slupska, Julia, Jeanette Lowrie, Lilly Irani, and Deian Stefan. "How Secrecy Leads to Bad Public Technology." (2021). https://escholarship.org/content/qt89d7826n/qt89d7826n.pdf

[xxviii] Amoore, Louise. 2020. Cloud Ethics: Algorithms and the Attributes of Ourselves and Others. Durham, NC: Duke University Press, 18-19.

[xxix] Dubal, Veena. "The Time Politics of Home-Based Digital Piecework." *Center for Ethics Journal: Perspectives on Ethics*, July 4, 2020. https://doi.org/10.2139/ssrn.3649270.

[xxx] Costanza-Chock, *Design Justice*; Hamraie, Aimi. "Designing Collective Access: A Feminist Disability Theory of Universal Design." *Disability Studies Quarterly* 33, no. 4 (2013).

[xxxi] Amoore, *Cloud Ethics*, 16.
[xxxii] Collier, Stephen J., and Andrew Lakoff. "Introduction: Systemic Risk." *Limn* 1, no. 1 (2011). https://escholarship.org/uc/item/0j88z94p.

[xxxiii] Felstiner, A. 2012. "The Weakness of Crowds." *Limn*.

[xxxiv] Roach, Levi. *Forgery and Memory at the End of the First Millennium*. Princeton University Press, 2021.

[xxxv] Sunder Rajan, Kaushik. 2006. *Biocapital: The Constitution of Post-Genomic Life*. Durham, NC: Duke University Press, 115.

[xxxvi] Eubanks, Virginia. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press, 2018.