

When a Bot Scores your Karma: Algorithmic Ranking Systems as Uncertainty Reducers in Platform Gig Work

Katharina Lix and Melissa A. Valentine

Stanford University

{klix, mav}@stanford.edu

Abstract

Drawing on a multimethod study of a network-based platform company, this article investigates the conditions under which freelancers develop trust and engagement in relationships mediated by algorithmic evaluations. In the network we studied, freelancers developed intense engagement focused on an algorithmic ranking system (referred to as “karma”). This engagement was reflected in prosocial behaviors and a willingness to work on the network’s behalf and dissipated when the system was discontinued. This collective engagement contrasted to freelancers’ day-to-day reactions to the system, which were more consistent with prior literature characterizing freelancers’ perception of algorithmic systems as opaque and coercive. Our findings highlight three related mechanisms underlying freelancers’ engagement with the algorithmic rankings. First, the algorithmic rankings were used in a way that reduced uncertainty freelancers faced about critical processes. Second, the algorithmic rankings were calculated in ways that increased perceptions of procedural justice -- not because the scores were perceived as accurate, but because they were consistent and documented in software code; perceived procedural justice in turn supported trust and engagement. Third, the rankings produced shared emotional experiences, including gamification dynamics and collective action. These inductively theorized mechanisms contribute new understanding about employment relationships mediated by algorithmic systems.

Keywords: algorithms at work, evaluations, ranking, networks, platforms, engagement

Work groups are only viable when their members are willing to engage in discretionary cooperative behaviors: those voluntary acts of cooperation that are not required or defined by formal rules or job descriptions, but that help facilitate the group's social and functional interactions (Tyler and Blader, 2000: p. 4; Tyler and Blader, 2003). Researchers have explored various conditions related to high levels of group engagement (see e.g., Wollard and Shuck (2011) for a review), but many workplaces are currently introducing novel employment and managerial relationships that have unknown consequences for group engagement. Of particular relevance is the increasing reliance on software systems that in various forms -- for example, platforms, bots, and online labor markets -- are used to accomplish managerial functions of hiring, monitoring, encoding performance feedback, disciplining, and firing (Kellogg et al., 2020). New research is needed to explore how these changes to the cultural schemas of systems of organizing -- described variously as a shift from paper-based bureaucracy to "algocracy" (Aneesh, 2006), or from organizational charts to "organizational operating systems" (Valentine and Hinds, 2020; Bahat, 2015) -- are influencing members' algorithmically-mediated relationships with their work groups and employers.

Prior research offers some insight into the question of how the shift toward algorithmically mediated relationships affects employees' engagement with the organization, but also some unclear predictions. The antecedents of engagement, reflected in employees' willingness to invest themselves on their group's behalf (Tyler and Blader, 2003; Rothbard, 2001), have received less attention from researchers than its benefits (Shuck and Wollard, 2010; Wollard and Shuck, 2011). The consensus in prior research seems to be that engagement depends, among other factors, on high levels of trust among group members (Schaufeli et al., 2008; Shuck et al., 2011; Herzberg, 1959, 1968; Tyler and Blader, 2003). Critically, members are most likely to develop trust and

subsequent engagement when they believe that interpersonal treatment and decision-making in the organization are based on principles and processes that are fundamentally fair (Tyler and Blader, 2003; Colquitt et al., 2012). Also, members tend to develop trust and engagement with groups and organizations that grant autonomy and then recognize, value, and reward discretionary cooperative acts (Brown and Leigh, 1996; Shuck et al., 2011).

What remains less clear is whether and how group members develop trust and engagement in work relationships that are shaped and mediated by algorithmic software systems. The research literature focused specifically on characterizing algorithmic work systems describes these systems as likely opaque and occasionally alienating because the system user-interface can obscure the kind of joint sensemaking of different situations that might result in a group or manager allowing discretion or autonomy (e.g., Thompson and Vincent, 2010; Wood et al., 2019; Chan and Anteby, 2018). The related studies that have examined specifically how people respond to these systems have tended to find resistance and gaming of algorithmic systems, and frequent disengagement from the platforms (e.g., Bolton, 2004; Gill et al., 2019; Hodgson, 2004; Lipsky, 2010). A smaller number of studies have found high levels of interaction – almost obsessive involvement – with algorithmic systems that “gamify” desired work behaviors and outcomes (Deterding et al., 2011; Edery and Mollick, 2009; Mollick and Rothbard, 2014; Petre, 2018; Walz and Deterding, 2015). Some scholars have argued, however, that gamification is an exploitative control system that ultimately alienates workers and leads them to disengage from the organization (Bogost, 2015). More research is needed to develop a clearer understanding of the conditions under which trust and engagement develop in work relationships mediated by algorithmic systems.

In this paper, we aim to contribute to this nascent research literature by analyzing an extreme case where a network of gig contract employees developed high levels of trust and

engagement with their work group, while many interactions and decisions were mediated by an algorithmic ranking evaluation system. We analyzed data from our 18-month field study of a gig platform company that organized a network of over 800 freelance “gig” software developers into teams that worked on on-demand software development projects for individual and corporate clients. The platform used an algorithmic evaluation system referred to as “karma” to publicly rank freelancers and to make staffing decisions. The system assigned each freelancer a “karma score” to reflect automatically collected, real-time data about their project performance and feedback they received from other freelancers. Both qualitative and quantitative analyses revealed that even though in their day-to-day interactions, freelancers often perceived the system as coercive and flawed, and engaged in resistance tactics consistent with findings from prior literature, they also developed high levels of engagement focused around karma. This engagement was reflected in frequent prosocial behaviors and a heightened willingness to work on the network’s behalf. Our data revealed a significant drop in engagement when the algorithmic ranking system was discontinued and replaced with a non-algorithmic evaluation system that did not include a public ranking.

We draw on the theory of group engagement and an inductive analysis of our data to theorize the mechanisms underlying freelancers’ engagement in response to the karma system. First, the network-based platform model was rife with uncertainty about accessing jobs, forming relationships, and gauging individuals’ status and performance. Freelancers used the algorithmic ranking as a strong “uncertainty reducer” about these critical processes. Second, the algorithmic rankings were calculated and used in ways that increased perceptions of procedural justice -- not because the scores were perceived as accurate reflections of individuals’ quality, but because they were consistent and documented in software code. Perceptions of procedural justice supported

trust and engagement. Third, the rankings were configured in ways that produced strong and shared emotional experiences including gamification dynamics and solidarity. These inductively theorized mechanisms contribute new understanding about employment relationships mediated by algorithmic systems. In sum, this paper contributes new insights about how algorithmic evaluation systems can become a means through which distributed members navigate the complexities and uncertainties of network-based employment.

GROUP ENGAGEMENT

Group members' level of engagement -- defined as members' focus of attention on and willingness to invest themselves on the group's behalf (Rothbard, 2001; Tyler and Blader, 2003) -- relates to critical group outcomes. In organizations, employee engagement has been linked to reduced turnover, retention, enhanced job performance, and organizational citizenship behaviors (Maslach et al., 2001; Saks, 2006; Shuck et al., 2011). Engagement has also been associated with productivity, discretionary effort, affective commitment, levels of psychological climate, and customer service (Christian et al., 2011; Fleming and Asplund, 2007; Rich et al., 2010; Richman, 2006). In addition, employee engagement appears to be predictive of firm financial performance (Ott, 2007) and growth (Xanthopoulou et al., 2009). From the literature, it is clear that members' level of engagement with their group or organization matters.

Engagement is reflected in group members' willing discretionary cooperative behaviors - that is, in voluntary acts of cooperation that are not required by the group's formal rules or one's job description (Tyler and Blader, 2000: p. 4; Tyler and Blader, 2003). Examples of discretionary cooperation behaviors include helping a colleague in need or taking on other voluntary tasks to improve the group's overall functioning and sense of belonging (for example, picking up one's trash from the office floor (Tyler and Blader, 2000: p. 4)). They may also include other social

behaviors that indicate members' emotional and cognitive focus on the group and its outcomes, such as checking in, offering help, and sharing personal information to build a sense of relatedness and community (Wollard and Shuck, 2011; Rothbard, 2001).

While prior research has identified a multitude of potential antecedents of group engagement, members' trust in the organization appears to be one of the most important factors for its emergence. Trust is defined as "confident, positive expectations about the words, actions, and decisions of another in situations entailing risk" (Colquitt et al., 2012: p. 1). Critically, employees are most likely to develop trust when they believe that decision-making and interpersonal treatment in the organization are based on principles and processes that are fundamentally just (Tyler and Blader, 2003). Related research shows that employees also need to also feel *trusted* by their organization to become engaged with it (e.g., Edwards, 1979; Jaros, 2010; Thompson and Van den Broek, 2010). Feeling trusted by the organization involves an organizational culture in which employees feel empowered to act on their own and in ways that go beyond their job description, and where their discretionary cooperative acts are explicitly recognized and valued (Brown and Leigh, 1996; Shuck et al., 2011).

CHANGING EMPLOYMENT RELATIONS AND ALGORITHMICALLY-MEDIATED WORK RELATIONSHIPS

The theory of group engagement was developed as a general framing of the social psychology of human groups (Tyler and Blader, 2003); much of the research literature developing and testing insights related to this theory in work groups and work relationships has focused on traditional bureaucratic organizations, where managers and members may have long-standing relationships over time, and where the organization as a workplace is a strongly institutionalized setting with clear norms, established onboarding and socialization processes, and hiring processes

that match hires directly to well-defined tasks nested in well-defined jobs in specific work groups (e.g., Saks, 2006; Macey and Schneider, 2008; Shuck et al., 2011; Penner et al., 2005). Yet as many researchers have noted, an increasing number of work arrangements are not characterized by these kinds of relationships with managers and work groups (e.g., Christin, 2018; Phanish et al., 2014; Valentine et al., 2017; Orlikowski and Scott, 2014). Instead, companies and organizations are pursuing two key interconnected socio-technical changes that are reconfiguring employment and work relationships.

First, many companies are developing networks and platforms that are used to complete knowledge-intensive tasks for internal or external clients; these configurations are sometimes referred to as network-based models of organizing (e.g., Phanish et al., 2014; Horton, 2010; Boudreau & Lakhani, 2013; McAfee and Brynjolffson, 2017). In network-based organizations, a company will develop and curate a network or crowd of expert individuals who are staffed to teams and projects dynamically and in an on-demand fashion. The members of these networks are freelancers or “gig workers.” They are rarely full-time employees with the career trajectories and employee benefits typical of full-time employment. Network-based organizing takes various forms (McAfee and Brynjolffson, 2017; Phanish et al., 2014). Examples include “flash teams,” where clients convene teams from the Upwork labor market (Valentine et al., 2017), [Deloitte “Pixel,”](#) where Deloitte curates crowds and then staffs members to new product development teams, [Artella](#), which organizes a network of animators to dynamically staff animation projects, and [b12](#), which curates a crowd of expert content creators and web developers for website development. Other examples of network-based organizations focused on composing teams of pre-vetted engineers and data scientists to work on on-demand client projects include [Factored.ai](#), [MVP Factory](#), [Terminal](#), [TopCoder](#), [TopTal](#), [Turing](#), [Distributed](#), and [Comet](#). Given the recency of

network-based organizations' emergence, relatively little research has examined individual workers' experiences in them. However, our observation in learning about this way of organizing and in conducting the present study is that people working in the network-based model confront different rules, relationships, complexities, and uncertainties than those in traditional organizations.

Companies shifting toward a network-based model are confronted with new management challenges (e.g., Fayard et al., 2016): Not only must they organize large groups of highly distributed and diverse workers. Oftentimes, they cannot rely on stable reporting relationships between managers and workers to monitor performance, set incentives, and ensure retention. These challenges give rise to the second socio-technical change that many companies are pursuing: The increasing use of data and data analytics, including machine learning, to analyze information and make decisions - especially about work processes such as staffing and evaluation (e.g., Faraj et al., 2018; Kellogg et al., 2020; Bailey et al., 2019; Karunakaran, 2016, Kittur et al., 2019; Lehdonvirta et al., 2019, Leonardi and Contractor, 2018; Ramamurthy et al., 2015; Rahman, 2018; Rosenblat, 2018; Varshney et al., 2014). In addition, more and more organizations use "algorithmic evaluation" to aggregate diverse data on employees' behavior and performance, and combine it with algorithmically or human-generated evaluations of them from within or outside the organization, for the purpose of quantifying and ranking employee performance (e.g., Barrett et al., 2016; Christin, 2018; Horesh et al., 2016).

At least in theory, using data, data analytics, and algorithms to augment or replace the role of human managers has two key advantages: First, it allows for managing distributed work forces at larger scales, lower cost, and faster speed. Theoretically, there is no limit to the number of workers an algorithm programmed for this purpose can oversee and make near instantaneous

decisions about. Second, algorithms are often believed to make more accurate and less biased decisions than their human counterparts. Even though recent research has cast a spotlight on prevalent problems of algorithmic bias (e.g., O’Neil, 2016), the necessity of human curation to make most algorithms function correctly (e.g., Pine et al., 2016), and the profound effect the mere presence of algorithmic technologies can have on group collaboration and coordination outcomes (Beane and Orlikowski, 2015), algorithms designed to perform managerial tasks are met with rising enthusiasm by leaders across private and public organizations (Kellogg et al., 2020).

Given the purported benefits of algorithmic management systems, algorithmic evaluations and ranking systems are especially prevalent in gig labor platforms and network-based platform companies like those characterized above (Rahman, 2018; Cameron, 2020). They allow organizations to quickly sort members of large and distributed workforces for the purposes of allocating tasks and rewards (Cameron, 2020). The nascent but growing literature on algorithms in the workplace suggests that people meet these systems with gaming or resistance when they create information asymmetries between employees and their employers that favor the latter, and when they are perceived as opaque, unjust, or condescending (Kellogg et al., 2020; Valentine and Hinds, 2019; Mollick and Rothbard, 2014; Christin, 2017; Thompson and Vincent, 2010). Additionally, being monitored through ongoing data collection or interactions with bots or apps can signal to employees that the organization does not trust them, and can reduce employees’ trust in the organization (Anteby and Chan, 2018; Boyd and Crawford, 2012; Eubanks, 2018; Noble, 2018; O’Neil, 2016; Pasquale, 2015; Scholz, 2012; Zuboff, 2019). Many of these prior studies focused on the use of algorithmic management systems in independent jobs such as driving for ridesharing platforms (Cameron, 2020), image labeling (Gray and Suri, 2019), or solo freelance contracting (Rahman, 2018). To our knowledge, no study has yet characterized algorithmic

management and its consequences in a network-based organizing model where expert freelancers were interdependent and collaborative in their work. Research is needed to explore group dynamics in this setting, particularly because of the previously identified problematic effects of algorithmic management systems on trust.

In this paper, we aim to develop new theory for this nascent research literature by analyzing an extreme case where a network of over 800 gig contract employees developed high levels of trust and engagement in their relationship with a platform company and work group. Their relationships were mediated by an algorithmic ranking evaluation system dubbed “karma.” We analyzed data from our 18-month field study of this gig platform company. Both qualitative and quantitative analyses revealed that even though freelancers felt negatively about the system’s coerciveness and context-ignorance, they developed high levels of engagement that focused around “karma.” Our data revealed a significant drop in engagement when the algorithmic ranking system was discontinued. We integrate an inductive analysis of our data with the group engagement and uncertainty management literatures to demonstrate how members enacted and interpreted the algorithmic ranking system as a strong “uncertainty reducer” in a work system that was otherwise strongly characterized by precariousness and ambivalence.

METHOD

Research Site

Our aim was to develop new theory by conducting multimethod field research (Edmondson and McManus, 2007). We studied a network-based organization, “FlashCo” (pseudonym), that composes entirely remote teams of highly skilled freelancers to work on on-demand software development projects for individual and corporate clients. During our 18-month observation period, about half of the freelancers on FlashCo were full-time freelancers, while others had full-

time jobs in other organizations. Some were employed at technology companies, and a considerable number used FlashCo as a source of supplemental income while working on bootstrapping their own startups. FlashCo provided only a minimal structure within which teams operated. Its main activities lied in vetting freelancers before they were admitted into the network by way of a technical interview process, in the recruitment of clients, and in the matching of freelancers to project teams. The teams were managed and led by project managers (PMs), who were also freelancers in FlashCo's network. During the time of this research, FlashCo headquarters staff ranged from 30 to 50 employees, and the network of freelancers ranged from 200 to 800 engineers, designers, and project managers. FlashCo received funding from venture capital, which gave them time to develop the network and the platform (similar to dynamics described in Shestakofsky (2017)).

Karma System for Algorithmically Evaluating FlashCo Freelancers

Part of the FlashCo leadership team's plan for developing the network and automating and scaling the processes used for hiring, staffing, and managing the projects was to create an algorithmic ranking system for freelancers that they referred to as "karma." Every person in the FlashCo network (not including headquarters staff) received a karma score when they joined the network. This initial score was based on their interview performance, but the variance in initial karma score assignments was relatively small. This karma score was publicly displayed on freelancers' Slack profiles.

Freelancers had two avenues for increasing their karma score. One, delivering milestones on time and achieving high customer satisfaction scores as part of a project team earned each team member a constant amount of karma points, which was specified in an internal Wiki that all

freelancers had access to. Conversely, when teams failed to deliver milestones on time or received low customer satisfaction ratings, each member lost a constant amount of karma. Slack bots, automatically programmed agents designed to interact with other chat users via conversation, were configured as the interface between the freelancers, this database, and the software that calculated the algorithmic ranking scores. The Slack bots appeared in a project teams' chat channel and announced karma awards or deductions in real time, but they were not configured to process or respond to team members' chat messages. Since team members could not meaningfully interact with the bots, they represented a one-way communication channel.

Second, network members could receive small amounts of so-called "peer karma" from other members of the network as tokens of appreciation or gratitude. New network members sought to earn peer karma by performing favors or so-called "chores" - well-defined sub-tasks from existing projects, such as vetting a specification document, conducting a code review, or doing research - for other members of the network in exchange for modest amounts of karma. Peer karma awards were not subtracted from the karma balance of the individual giving the award. Freelancers could also request karma subtractions for peers who they believed had shown transgressive behavior using the same procedure, although this was much less common.

The number of karma points granted or subtracted automatically by the algorithm was significantly greater than the number of points that freelancers could award or subtract each other through peer karma. The automatic karma adjustments came from teams' milestone timeliness and customer satisfaction ratings. The size of peer karma awards was constrained by headquarters. Thus, even though peer evaluations factored into karma scores to some extent, karma scores were largely determined by the algorithm that executed automated karma decisions based on teams' trace data. Freelancers in the network were aware of this difference.

The software code that computed individuals' karma scores based on their teams' project performance and peer feedback was written by one of the early headquarter staff members. Everyone at headquarters and some freelancers in the network knew who had authored the code and had a sense that if they wanted to look at the code, they could. The policies governing karma awards and deductions were also specified in an internal Wiki that was maintained by headquarters and that all members of the network had access to. Moreover, freelancers could view the "karma leaderboard," a public ranking of all network members based on their karma scores, by logging into the platform's online system that was used for processing payments, submitting peer karma requests, and retrieving documentation. All network members who were active on projects were required to log into this system on a daily basis for work, and therefore saw the leaderboards frequently. The data collected on milestone timeliness, customer feedback, and peer and PM feedback were stored in a database that only a few headquarters staff could access. It was subject to common database issues involving lagged data updates and occasionally erroneous entries. The karma algorithm thus sometimes made decisions based on erroneous data or miscalculations, but freelancers were not able to see the data or computations that had led to specific karma decisions.

Data

We collected and analyzed qualitative and quantitative data related to the karma algorithmic ranking system during our 18-month field study of FlashCo. The first author engaged in participant observation during this time, collaborating on a project at headquarters as an unpaid analyst. She helped a headquarters team develop a novel platform feature and used this unpaid analyst role to develop a close understanding of the FlashCo network and platform. Over the 18 months, she attended meetings at headquarters, contributed to ongoing projects, and participated in Slack channels, spending about 10 hours per week on average in this analyst role. The first

author wrote field notes and memos on this participant observation. We conducted thirty additional interviews using a semi-structured interview protocol focused specifically on developing inductive understanding of the karma ranking system. The interviews ranged from thirty minutes to four and a half hours and were distributed across FlashCo roles, including freelance project managers, engineers/developers, designers, and headquarter staff members, including the leadership team. Four interviewees had been freelancers at FlashCo and after some time accepted part-time or full-time positions at headquarters. While only about three-quarters of interviewees were still active on FlashCo projects at the time of the interview, all of them had been active on projects for the period during which karma was in place and the months immediately following its discontinuation. The interview protocol followed accepted practices for conducting research interviews, including asking for stories of concrete experiences and events before asking for interpretations (Spradley, 1979). We judged that we had theoretical saturation when we were no longer hearing new or diverse patterns in respondents' stories and interpretations about karma (Glaser and Strauss, 1967). We heard little substantive divergence in accountings of karma; in a general sense, the stories consistently described highly emotional day-to-day reactions to karma that ranged from frustration, especially when the Slack bots announced point deductions, to elation in response to karma gains and the system's gamified aspects; and also a shared sense that there was an almost fervent focus on karma as a revered FlashCo institution that was integral to the community's functioning. Interviews were recorded and transcribed, with participants' consent.

We triangulated our extensive observational and interview data with quantitative and qualitative archival data (Jick, 1979). We negotiated access to two main archival data sources, both of which had to be anonymized before we could download them, per our research agreement with FlashCo. The first was an archive of the full team Slack transcripts for 117 teams collaborating on

software development projects during and immediately after the period during which karma was in place. These transcripts included data from 421 unique individuals (36% female), with an average team size of eight, and amounted to about 800,000 messages total.¹ On average, teams' Slack transcripts comprised 1,200 messages, and ranged from 1000 to 6000 lines when exported to a word processing document format. The second data source was an archive of a public Slack channel that freelancers used to discuss topics of general interest to the network. This channel was started shortly after the karma system had been launched, was actively used by freelancers while karma was in place, with activity rates declining steadily after karma's discontinuation. The channel included 18,500 messages sent by 53 unique individuals, corresponding to 30,000 lines when exported to a word processing document format.

Analysis

The qualitative data analysis followed an inductive theory development process (Charmaz, 2008; Glaser and Strauss, 1967). Memos and interview transcripts were read several times, resulting in records of high-level themes. We started to see high-level patterns of charged and emotional language describing people's engagement with karma. We also saw that the freelancers were using karma in different ways and for different purposes. Using these high-level themes as a broad framework, the first author then conducted line-by-line analysis of the interview transcripts (Miles and Huberman, 1984; Strauss and Corbin, 1990). This open-coding began with reading a thought unit (Charmaz, 2008) and asking, "what is this an example of?" and more specifically, "how is karma being enacted or used in this story?" This first analysis produced codes related to the freelancers' interpretations and use of karma. They included "using karma to build

¹ We did not have access to direct messages exchanged between team members, only to messages that team members exchanged in team-wide channels.

relationships with project managers who could allocate you good work”; “using karma to find mentors”; “interpreting karma as reflection of one’s own quality”; “interpreting karma as encoding ideals of meritocracy”; “interpreting karma as more important than money”; “using karma to navigate the network’s social landscape”; “feeling energized by one’s karma score,” among many others. These codes were intended to be exhaustive, documenting all themes raised in the interviews. A second round of coding focused on analyzing how it came to be that people perceived karma to be procedurally just and how it came to be that people were emotionally engaged with karma. These analyses were focused on understanding the relationships, and were informed by iterations with the research literature on group engagement, uncertainty management, and algorithms at work.

The quantitative data analysis began with a similar open-coding process that involved reading select team and public transcripts and coding them for high-level themes and patterns. Given the massive amount of text data we collected, open-coding all of these transcripts was not realistic. We did a comprehensive open-coding on ten of the team Slack channels. We saw that 62 of the 117 teams explicitly discussed karma in their project teams, with 132,715 messages mentioning karma. For the public channel’s transcript, we did one full read-through (500 pages) and open-coded each individual Slack message, with a particular focus on conversations evolving around karma. We also used these channels as a quantitative triangulation on people’s descriptions of community engagement before and after karma was discontinued. We calculated the number of messages sent per month over time and the number of unique posters per month over time. We also calculated and visualized several metrics related to freelancers’ engagement in the public Slack channel, such as seeking help, offering help, sharing personal information, and asking questions. These analyses are visualized in Figure 1 and Figure 2, respectively.

We use these data and analyses in the following ways. First, drawing primarily on interview reports from freelancers and headquarters staff, we describe what working in the network-based organization was like, with a particular focus on the unique uncertainties freelancers faced in this novel organizational form. Second, drawing on data from interviews and teams' Slack transcripts, we show how freelancers used the karma system to interpret and navigate these uncertainties. We show how the algorithmic properties of the karma system engendered resistance and frustration in freelancers' day-to-day interactions with the system, but at the same time improved their perceptions of justice in the organization, which ultimately led to the formation of trust and engagement. Further, we present evidence that strong shared positive emotional experiences engendered by the system's gamification and opportunities for collective action led freelancers to discount their critiques of the system and develop predominantly positive views of it. Third, drawing on all data sources, we present evidence that the karma system, through its function as an uncertainty-reducer and amplifier of positive emotional entrainment, boosted engagement and prosocial behavior in the organization, and that these behaviors decreased sharply when the karma system was discontinued and replaced by a non-algorithmic evaluation system that did not include a public ranking.

FINDINGS

The FlashCo freelancers developed active engagement in their network relationships that were mediated by the karma algorithmic evaluation system. Even though they were sometimes frustrated at the karma system's interventions, they drew on karma as a resource that helped them navigate relationships that were otherwise characterized by uncertainty. They interpreted the system as transparent and consistent, thereby enhancing perceptions of procedural justice. They also felt and expressed shared emotions as they collectively interacted with the system, which was

configured to provide consistent and immediate feedback. The freelancers' engagement decreased sharply when the karma system was discontinued and replaced by a non-algorithmic evaluation system that did not include a public ranking.

Freelancers Confronted Considerable Uncertainty in the Network

FlashCo freelancers experienced considerable uncertainty when joining and working in the network. Their descriptions of navigating this uncertain landscape matched a sort of maverick mentality of prioritizing personal freedom in the "freelance lifestyle" at the accepted cost of having to figure out how to thrive in the "wild west" of the FlashCo network. They tended to contrast their freelance experience with their traditional work experiences, and said that the uncertain FlashCo landscape required them to spend considerable time and energy to figure out how to navigate and adapt with respect to 1) accessing jobs, 2) forming social relationships, and 3) gauging individuals' status in the organization. Table 1 contains illustrative quotes from freelancers' interview reports and Slack conversations demonstrating the sources of uncertainty they faced.

Starting Up and Accessing Jobs

Even though entering the network required passing a series of technical interviews, network members were not automatically guaranteed access to jobs. Headquarters assigned each new member an initial karma score based on their interview performance and gave them access to the network's internal communication tools and Wiki. Other than that, network members were "*on their own*" (freelance project manager 1, in an interview) to secure their first role on a project. The network's internal Slack channels functioned as job boards and discussion forums, where project managers advertised open roles on projects, and members exchanged information of general interest to the community, such as trends in the software development industry. To be staffed on

jobs, freelancers reacted to job postings in public Slack channels by sending a short message to the project manager that included their karma score. For most new freelancers, this presented a challenge, since they first had to earn additional karma points on top of their initial karma score in order to become eligible for roles. One freelance developer described the uncertainty she faced seeking her first role on a project as follows:

“So, you get on Slack, you’re in these channels with a few hundred strangers, and you don’t know anyone. What do you do? You need to get your karma score up quickly, because if other people see that your karma hasn’t increased since you joined, they’re not going to hire you. So, you gotta start talking to people, strangers, on Slack. You don’t know them, but you’ve got to find projects or jobs that will get you karma. But you’re flying blind in those first few days and weeks.” (Freelance developer 1, in an interview)

Forming Relationships in the “Wild West”

Navigating this “*wild west*,” as freelancers described it, meant figuring out how to form relationships and follow social norms in an eight-hundred-person distributed network that relied exclusively on online text-based communication. Newcomers knew little about other individuals they encountered in public Slack channels, except for the information these individuals shared in online discussions and in their online chat profiles, such as their name, professional role, karma score, and sometimes profile photo. New members could rarely rely on information about shared interests, nonverbal cues, or happenstance interactions to start conversations. There were few explicit norms and rules, but those who failed to read the “code” and violated the community’s norms had difficulty integrating. On project teams, these norms included fast response times, an upbeat tone, and a high degree of flexibility in terms of working hours and processes.

Gauging Relative Status and Performance

A final source of uncertainty that freelancers described was their own and others' relative status in the network. Because they competed against each other for access to lucrative projects, they wanted to know their own performance compared to other network members. Otherwise they were not able to estimate the kinds of jobs and associated salary ranges that they could reasonably expect to access. Additionally, freelancers wanted to be able to assess the credibility of individual contributors as they tried to make sense of the large volumes of rapidly evolving information posted in public channels, not least because they desired to gauge whose opinions on network-internal policies were likely to hold sway with decision-makers at headquarters. Many freelancers also hoped to learn from more experienced software professionals, but because they lacked visibility into the quality of work that others who were not on their team delivered, identifying the most skilled individuals in this large network was challenging.

----- INSERT TABLE 1 ABOUT HERE -----

Freelancers Used Karma to Navigate the Uncertain Landscape

Freelancers used the karma scores and karma rankings to make sense of and navigate these uncertainties. They used karma scores to understand how to access jobs, form and evaluate relationships, and gauge their own and others' relative status and performance. Table 2 displays additional evidence illustrating how freelancers used karma to navigate these uncertainties.

Accessing Jobs

Freelancers learned that jobs were allocated by freelance project managers, who used karma scores to decide which freelancer to choose for a specific role. Freelancers felt they had clear goals to aim for as they sought to secure jobs of certain durations and salary ranges based on

the minimum karma requirements for each job and project managers' adherence to karma in staffing decisions. A freelance developer describes the sense of predictability and transparency that karma provided on the decision-making process:

"It was very clear [how to get good work]: you get karma, you get gigs. [...] there was a clear path forward for you. [...] Now [working full-time at a technology company], I don't have that. I don't really know how management decides who gets to work on what." (Freelance developer 1, in an interview)

Another freelance developer said understanding karma's role in staffing decisions helped reduce the uncertainty freelancers faced with respect to their future income and career prospects:

"The sense of transparency [from karma] was really good ... They called it out how they chose their dev[eloper]s. That's really important if you're a freelancer because you're looking for as continuous an income as you can possibly get. That part felt good. If you wanted to get into a good project, you had a target to aim for." (Freelance developer 2, in an interview)

Forming Relationships

Karma also reduced uncertainty about forming relationships in the network. Relationships were important because the freelancers worked interdependently on project teams and influenced each other's job prospects through awarding and subtracting each other karma points.

Bonding over karma. Karma provided freelancers with a shared frame of cultural reference for forming relationships. Karma was a safe topic of conversation for freelancers, most of whom were strangers, because *"everyone had something to say about it"* (freelance project manager 1). Freelancers viewed karma as an institution that was at once *"wonderful and weird"* and unique to FlashCo. Humorous comments about karma were a common way for freelancers to *"break the ice"* when they first encountered their team members in an online chat at the start of a new project. Freelance project manager 9 described how teams started conversations by *"teasing*

whoever was on top of the leaderboard that day.” The casual conversations such comments sparked often lasted throughout the day and ran in parallel to teams’ technically focused exchanges. Karma was a versatile conversation topic; its meaning was shared and its relevance clear.

The freelancers also bonded over karma when explicitly forming and sensemaking their relationships. The following representative example from a conversation between two team members illustrates this. An engineer broke a karma threshold that made him eligible for highly lucrative projects on the platform. The project manager, who himself has a high karma score, addressed the engineer:

Freelance project manager 2, in a team Slack channel: *“You[‘re] now an Engineer with 330+ karma, the whole universe [is] in front of you, please don’t fuck it up [emoji] [emoji] ... Whenever you need help or a job in the future, just let me know and I can hook you up.”*

Freelance developer 3: *“Amazing, buddy! Thanks!”*

The project manager and the engineer formed an alliance rooted in mutual respect and reciprocity. Through their karma scores, they explicitly recognized each other as capable and committed peers.

Learning and shaping norms. Karma also allowed the freelancers to learn about and shape social norms. Everyone was united in their desire to earn karma points, so performing favors for others – even if they were strangers – in exchange for points was common. For example, freelancers readily asked each other questions and helped each other, even if they had no pre-existing relationship. Freelancers posted specific technical questions in public Slack channels. Even with no formal requirement to reward answers with karma points, doing so was the norm. As freelance developer 1 described in an interview:

“Sometimes I needed specific [help]. I needed someone to pair program with me because I [was] real stuck. So, if you found someone to help you with your problem, you would give them karma. I did that a bunch of times ... It was definitely an incentive for people to help each other.”

Freelancers embraced small karma awards as a means to express gratitude and appreciation for others’ help. A member of headquarters staff, who was in charge of reviewing and arbitrating karma requests, described how contractors used karma requests to encourage prosocial behaviors:

“Freelancers awarded each other karma a lot. A lot of [karma requests] were really cute. They could be the littlest things, such as, ‘This person helped me answer this random personal question’ or ‘This person made my day - [please give her] plus one [karma point].’”

Similarly, freelancers used karma deductions to penalize behaviors they felt were counterproductive. For example, freelance project manager 3 described how she used karma to sanction behavior that she considered inappropriate on her teams:

“I deducted karma when someone on my team was not responding. Also, when they used bad language because that destroys the climate. The first time someone used bad language, I would give them a warning. Second time, maybe another warning. But the third time, I would have karma points taken away. Then they would stop and course-correct.”

These and other reports from informants suggest that peer-to-peer karma awards and deductions were effective at maintaining and even strengthening the network’s social norms. Over time, the karma system contributed to a culture where prosocial behaviors, including helping one’s team members and other freelancers on the network, were the norm. Freelance developer 4 summarized in an interview:

“The karma system made some people act as if they were more team players than they actually were. The result was good. People across the whole network were just working more as a team.”

Finding mentors. Another way that karma encoded norms was to visibly identify individuals who succeeded on FlashCo, and those leaders became role models. Others watched their behaviors closely in order to learn the network's norms. These leaders felt a strong sense of pride in their karma score and felt responsible to help younger members of the network learn the "*FlashCo way*." Freelance project manager 4, a senior network member, summarized:

"As a new member, I can [use karma to] find someone and say, like, 'Oh, you're clearly more experienced than me and I need help.' And as an experienced member, I can also clearly see, 'Oh, maybe this person is doing something improper, so to speak, because they're a new member, or they have less experience, and so rather than me being a dick about it to them, I can help them learn the FlashCo way.' This was a common thing."

Karma also helped freelancers form relationships by signaling individuals' viability as mentors and future collaborators. Over half of our informants reported actively reaching out to individuals who held top positions in the karma leaderboards to ask for career advice or coaching on technical questions. Freelance developer 5 described how he used karma as follows during an interview:

"I looked at karma as a tool [to answer the question], 'Who should I learn from that really understands how FlashCo works?' and, 'Who's a good person to potentially partner up with in the future for certain projects?'"

Generally, the individuals atop the karma leaderboards were receptive to junior members' attempts to reach out for mentorship and advice. Freelance project manager 1, who led the karma leaderboard for multiple months, recounted in an interview:

"I would have people write to me and say, 'Hey, so cool you're on top of the karma leaderboard. What are some of the pointers you can give me?' And I'd help them out. [...] So karma [was] a tool that people used to navigate FlashCo."

Gauging Individuals' Status and Performance

Relatedly, karma helped freelancers explicitly gauge individuals' relative status and to place themselves within the network's social hierarchy. For example, during online discussions

that involved many contributors, freelancers checked the karma scores of those with whom they had not interacted to get a sense for their credibility and relative standing in the network. Freelance developer 2 described in an interview how he relied on karma scores to size others up, even though he was aware that karma scores were merely a rough approximation for others' quality:

“If I saw someone say something that I thought was cool, I'd click on their [Slack] profile and check their karma. I had to do this all the time. I would check to see if they're good [...] But then if someone said something stupid, I would be like, 'Oh shit, look at his low karma. This guy doesn't know anything.' It was kind of lame, but it was something that I couldn't avoid.”

Individuals with high karma scores were respected and referred to as the “elite” by other freelancers in the network. The opinions they expressed in public Slack channels were taken seriously not only by other freelancers, but also by members of headquarters. In an interview, freelance developer 6 explained the social hierarchy that karma established:

“People respected you if you had a high karma score, even if they didn't know anything else about you [...] People were really proud of that because [a high karma score meant] you've been around for a while. [If you had a high karma score,] what you said on Slack mattered, people listened to you.”

Similarly, freelance developer 4 described how he used his teammates' karma scores to get a rough sense for their quality as professionals and teammates:

“If I found someone that had 500 [karma points], I would say, 'Yeah, this is a good baseline that I'm working with someone that's committed to FlashCo and to their projects and who is responsible.' ... But below 500, maybe this person has not been in the network enough or has not worked enough for it.”

----- INSERT TABLE 2 ABOUT HERE -----

Karma Supported Perceptions of Procedural Justice

As described above, Karma was central to organizational life in the network. The system reduced critical uncertainties that freelancers faced by facilitating relationships and by helping

them interpret the network's vast social landscape. Nonetheless, freelancers' awareness of the system's shortcomings was evident. The greatest source of frustration and resistance was that the karma algorithm did not take into account teams' context as it automatically awarded or deducted points based on the team data it collected for a given milestone. For example, in some cases, deliverables that were marked in the system as "late" had in fact been delivered on time by the team, but the client had forgotten to mark them as complete in the online system. In such cases, the karma algorithm automatically deducted karma points with limited appeal options.

When karma deductions based on a "false read" of the data occurred, freelancers' frustration was compounded by the incessant reminders and warnings that karma chatbots sent to them. Freelancers perceived these bots as coercive and felt powerless against them, especially when they believed that the karma decision had been unfair. On one team that had lost karma due to a systems error and consequently received three warning messages from a chatbot, a freelancer wrote to his team: *"This bot is driving me INSANE. It's impossible to concentrate with all the flashing and pinging. Isn't there a way to just turn it off?! I'm sick of this."* Teams' Slack conversations vividly showed freelancers' frustration at the bots' "terrorizing" and "stupidity," since the algorithms underlying them were context-unaware and unreceptive to explanations.

The karma system's shortcomings notwithstanding, freelancers fixated on karma as the primary metric for evaluating themselves and others. What is more, the very properties of the algorithm that sometimes frustrated freelancers also gave rise to perceptions of the system as highly consistent. Freelancers felt that karma was just in the sense that it made decisions based on clear rules and that it treated all freelancers the same. For example, because open roles were associated with minimum karma requirements, freelancers felt that the process underlying staffing decisions was transparent and fair. In interviews, freelancers highlighted that even though they did

not always agree with the karma system's decisions or its criteria for distributing rewards, they appreciated that the underlying algorithms applied the same set of rules to every member of the network. To many freelancers, this represented a key advantage over human managers, who, in the words of one freelance developer, "*can't be trusted because they're biased, they have favorites. They change how they make decisions without telling you why, and they don't treat everyone the same.*" Overall, a significant portion of freelancers seemed to favor being evaluated by the karma system over a potentially biased human manager.

Additionally, karma gave freelancers the impression that their opinions mattered in shaping the network and policing its norms. Because they were able to award and subtract each other karma points, freelancers felt that they could influence who succeeded in the network. They felt a sense of agency and voice in shaping the organization, even though the number of points freelancers were allowed to exchange was miniscule compared with the algorithm's karma awards and deductions, and therefore unlikely to significantly affect individuals' career outcomes in the network. However, when authorities give employees "voice" by letting them provide input into decisions, they perceive the resulting decisions as fairer than if they had not had the opportunity to make themselves heard (Thibaut and Walker, 1975; Leventhal, 1980). Karma improved freelancers' *perceptions* of the quality of decisions on the platform, even though their actual ability to influence these decisions was negligible.

In addition, freelancers appreciated that karma made visible acts of helping and "*going the extra mile*" that many said often go unnoticed in traditional organizational contexts. Freelance developer 1, who worked on FlashCo for multiple years and now holds a full-time job at a technology company, explained:

“[With karma,] your efforts were recognized when you helped someone out or when you collaborated. Whereas, without karma, that gets lost in translation [...] You don't get that [recognition that karma offered] at a real job. Just now [at my full-time job], I was pair-programming with somebody and nobody knows that I did that. But if there was karma, my boss would be like, ‘Oh, you got 10 karma. What did you get 10 karma for today?’ And I would say, ‘Oh, well, I was helping out another person on the team.’”

Echoing this informant’s argument, many noted that karma significantly contributed to the quality of interpersonal treatment in the network (see Table 3 for additional illustrative evidence). Karma provided a tangible means for incentivizing, rewarding, and eventually normalizing prosocial behavior. Freelancers checked in on each other, voluntarily helped each other troubleshoot issues, and provided each other emotional support during challenging times. Such extra-role behaviors have been shown to foster the emergence of positive affect, swift trust, and effective coordination in teams (Valentine, 2018). Moreover, when rules are perceived as consistent, as they were with the karma system, organization members perceive the organization as predictable and dependable, and are therefore more likely to trust it (Colquitt et al., 2012; Tyler and Lind, 1992). In this way, karma contributed to trust between freelancers and the organization. Trust, in turn, influenced members’ levels of engagement with their organization (Tyler and Blader, 2003).

----- INSERT TABLE 3 ABOUT HERE -----

Karma Produced Strong Shared Emotions

Karma improved freelancers’ perceptions of organizational justice with respect to decision-making and interpersonal treatment. In addition, karma supported interactivity and solidarity, which produced strong shared emotions among freelancers that further intensified their

engagement with the network. Table 4 shows illustrative evidence for the collective emotions freelancers experienced in relation to the karma system.

Gamification. The karma system relied heavily on gamification - the application of game-playing elements and principles. Many freelancers described the karma system's reliance on classical game elements, such as point scoring, competing with others, and rules of play, as "addictive." Some reported feeling "obsessed" with increasing their karma score. The feelings of rapture and elation many freelancers expressed in response to earning karma points were remarkable. Freelance project manager 5 recounted the entrainment that gamification produced as follows during an interview:

"I... liked that you got all of this reinforcement when you're working ... It feels like a game, and you're constantly rewarded as you do good things. Sometimes, I'd wake up and see that I got karma, and I'd think, 'What a great day! The FlashCo Gods were nice to me today!'"

Freelance developer 7, who reached a top ten spot on the leaderboard after a year and a half on the network, expressed reverence for the karma system as a "brilliant manipulation" that he willingly participated in. He set successively more challenging karma goals for himself and enjoyed tracking his progress toward them:

"I had some goals. Like, can I get to the top 50? Or maybe even top 20? And then to top ten? ... You had these Slack bots, which told you, 'You get plus three points of karma.' It was a great muse in my life because I could look at my phone at night and see that."

For some freelancers, the psychological rewards from engaging with the karma system were so strong that, in the words of freelance developer 8 during an interview, "earning karma became more important than earning money." Freelance project manager 4 described how some freelancers, including those in top positions on the karma leaderboard, were so eager to earn karma

that they would accept “*the completely irrational deal*” to work in exchange for karma points instead of money in order to earn “*bragging rights*,” adding, “*It was crazy.*”

The effects of karma’s gamification were evident also in intense collective emotions among freelancers. As freelance project manager 6 recounted, “*There was a moment of mania when [karma] was first introduced ... then, [the network] went through cycles of excitement and despair [about karma].*” On project teams, bots’ karma decisions affected all team members equally, regardless of their relative contributions to a particular performance outcome. As a result, team-wide karma events sent teams into collective states of ecstasy when they earned karma points (resulting in exclamations: “*We’re going to Vegas!*” (freelance designer 1, in a team Slack channel)), and into spirals of disappointment when they lost points (“*My life has no purpose.*” (freelance developer 9, in a team Slack channel)). In interviews, freelancers described how earning karma as a team made them feel closer and more connected to their teammates. Suffering karma losses together on a project gave freelancers a sense that they were “*in this crappy situation together*” (freelance developer 5, in an interview), which produced feelings of camaraderie and solidarity. In sum, the psychological effects of karma’s gamified elements contributed to the system’s uncertainty-reducing effects: Collective emotional experiences in response to karma produced a sense of relatedness and belonging among contractors.

Solidarity through collective action. Karma was culturally, socially, and economically important to the freelancer network. Many freelancers cared deeply about the policies that governed the allocation and distribution of karma points, and they developed a sense of solidarity when they collectively engaged with headquarters around karma policies. On occasion, network members banded together to force changes in response to new karma policies. In one case, to

express their outrage, freelancers instigated what one informant referred to as a “*Slack revolution.*”

Freelance project manager 6 described how it unfolded during an interview:

“Like most revolutions, it starts as hushed conversations in a Slack channel ... then it turns into a public discussion, and then it explodes ... [freelancers] rant[ed] and rave[d] on Slack, ... and then [the network organization’s founder] would have to come in and shut down the Slack channel.”

At other times, groups of contractors convened in private Slack channels to assemble formal lists of complaints, along with suggested corrective measures, which they then shared with headquarters in hopes of affecting policy changes.

With rare exceptions, freelancers’ collective actions were not successful at eliciting changes from headquarters. Nonetheless, in the words of one informant, the ongoing activism around karma amplified among freelancers a “*perhaps false impression that FlashCo was a community-owned platform*” (freelance project manager 1, in an interview). During periods of collective action, the karma system became a foil against which freelancers banded together. In advocating for changes to the karma system, they experienced a sense of collective purpose and agency. Freelancers formed and enacted a shared identity that was in part defined through opposition against the “*FlashCo overlords,*” as freelance developer 6 referred to headquarters. Freelancers’ level of investment in and protectiveness of the karma system produced a palpable tension in the network when freelancers felt that the system was under threat. Freelance developer 2, who viewed the network’s level of activism around karma more critically than most, described the atmosphere as follows during an interview:

“It was kind of ugly ... because ... you don't want to see people contesting the system that they're working under all the time. It feels unstable. It feels like a riot could break out any time.”

Overwhelmingly, however, freelancers recounted how the sense of community and their perceived ability to shape the network's norms and rules through banding together over karma decisions was chief among the reasons they preferred working on FlashCo over alternative forms of employment. Freelancers reported feeling a sense of control and ownership over "their" platform, which in turn reduced the salience of the uncertainties they confronted in the network. Moreover, the feelings of elation and solidarity that the karma system engendered helped freelancers cope with the frustrations they occasionally felt about the algorithm's ignorance of context and its lack of channels of appeal.

----- INSERT TABLE 4 ABOUT HERE -----

Changes in Engagement Levels due to Discontinuation of Karma

The karma system's importance for freelancers' levels of engagement with the platform was cast into sharp relief when the system was discontinued. Around February 2017, the karma system was taken offline for maintenance reasons, and was not put back online during the 12 months of our observation period following its discontinuation. The consequences in the network were far-reaching, as freelancers' levels of engagement with the platform plummeted.² In the absence of karma, many of the uncertainties that karma had helped reduce began dominating freelancers' experiences on the platform. Under the non-algorithmic system without a public ranking that replaced karma, freelancers felt upset at what they perceived as a lack of transparency into how decisions were made, what actions they needed to take to secure the jobs they desired,

² Our analyses do not allow for the identification of karma's discontinuation as the sole cause of the decrease in freelancers' engagement. However, nearly all informants highlighted the discontinuation of the karma as a main reason for their personal decision to participate in the network less and for the shift they observed in the network's overall culture.

and how they could plan for future income. These uncertainties persisted in spite of headquarters' announcements about basing staffing decisions on peer feedback and past performance. Freelance developer 2, who worked in the network while the karma system was in place and remained active for three years after it was discontinued, described:

“When karma was around, if you wanted to get into good projects, you had some [karma] target that you could aim for. And we kind of lost that after karma was gone. The process became more secret ... It became a bigger question mark for us: ‘How does the ranking system work [now that karma is gone]?’”

In addition, contractors lamented the loss of community on the platform, which they tied to the discontinuation of the karma system. Freelance developer 8 recounted the cultural changes resulting from the discontinuation of karma made FlashCo a less attractive place to work:

“[When karma was removed], I was like, ‘I know that [headquarters] had internal mechanisms to rank developers ... but there was nothing external-facing to do that, so for me [the network] really became a different place after that. I really wasn't as interested in working on FlashCo projects after that. ... I enjoyed [karma]. I was really shocked and very sad to see it go.’”

Freelance designer 2 echoed his sentiment in an interview:

“I feel like there's something that was lost [when karma was discontinued]. I'm not as active on FlashCo anymore.” ... “[When karma was around,] I knew a lot more about who else was in the network, what people were interested in, what people were working on and who I should be working with, who I shouldn't be working with, who was good, who was bad.”

Table 5 offers additional illustrative evidence about freelancers' decreased engagement with the network in response to the discontinuation of karma. This decline in engagement was also reflected in decreased levels of activity in a public Slack channel that freelancers used to discuss matters of general interest to the network and to ask each other questions or seek help. Figure 1, Panel A, shows the decline in the number of freelancers who were active in the public channel. Panel B shows the drop in the number of messages exchanged in this public channel.

----- INSERT TABLE 5 ABOUT HERE -----

----- INSERT FIGURE 1 ABOUT HERE -----

The cultural change associated with karma’s discontinuation was reflected not only in the changed frequency with which contractors engaged with the network’s social community, but also in the altered nature of the exchanges that occurred. As freelance developer 4 described in an interview:

“You lost the engagement ... Before, you would have more discussions, more people participating, [sharing thoughts such as:] ‘Here’s this startup. I find this technology interesting,’ et cetera. But now the Slack channels just feel so transactional.”

Multiple informants recalled a marked drop in prosocial and helping behaviors following the discontinuation of karma. Freelance developer 1, in an interview, said:

“There used to be so much more interaction on Slack. Conversations, people helping each other, [asking,] ‘Can you do this for me?’ It used to be a lot more than it is [now].”

Figure 2 illustrates the substantive differences in the messages exchanged in contractors’ public Slack channel between the time when karma was in place and when it was taken offline.

----- INSERT FIGURE 2 ABOUT HERE -----

DISCUSSION

This paper develops new theory to explain workers’ engagement in algorithmically mediated relationships with peers and employers. We examine the effects of an algorithmic evaluation and ranking system known as “karma” on engagement in an entirely distributed network of over 800 high-skilled freelancers collaborating as teams on software development projects. Our data reveal that the karma system facilitated engagement and prosocial behavior among organization members, and that engagement levels dissipated once the system was discontinued. Even though freelancers’ day-to-day interactions with the karma system were often characterized

by frustration and resistance in ways that are consistent with findings in the prior literature, the karma system contributed to engagement because it helped freelancers navigate the network's highly uncertain landscape and fostered perceptions of procedural justice. Freelancers developed trust in the karma system because its rule-based decisions were perceived as consistent, predictable, and unbiased, albeit not always as accurate. In addition, the system produced collective positive emotional entrainment through gamification and the elicitation of solidarity among freelancers, which further enhanced engagement levels.

These findings suggest that workers' responses to algorithmic control are contingent on the types and levels of uncertainty they experience and on their perceptions of procedural justice in the organization: When algorithmic systems meaningfully reduce uncertainties that workers face and provide a sense of consistency and predictability, they can support trust and engagement. This finding highlights the need for further in-depth studies of workers' engagement with and responses to algorithmic systems across a variety of organizational contexts. Our findings contribute to the algorithms at work literature, the theory of engagement in work groups, and research on rankings and evaluations.

Contributions to Algorithms at Work Research

Scholars of algorithms and organizations have argued that the "quantification of everything" (Espeland and Stevens, 2008; Christin, 2017), from employees' communication behaviors to their performance, is ushering major shifts in organizations with respect to expertise, occupational boundaries, organizational control, and coordination (Faraj et al., 2018). The dominant view in the literature is that algorithms create tight but hardly visible "iron cages" around workers (Faraj et al., 2018): Because these systems' rules are not as easy to understand or available

for scrutiny as decision-making processes in traditional organizations, workers are left at the mercy of powerful systems that control them (Kellogg et al., 2019). Prior research has shown that when workers perceive the algorithmic systems that manage them or that inform their actions and decisions as opaque, they can lose trust in their own cognition when these systems break down (Danaher, 2016) and begin to doubt their own decisions (Lebovitz et al., 2019).

Our study is resonant with this literature to some degree. The “black-boxedness” of karma elicited frustration, especially when freelancers believed that specific karma deductions were based on “false reads” of the data. As prior work predicts, freelancers in our study sought to reclaim control against the algorithm through resistance tactics that ranged from exploiting loopholes in the algorithm’s computational processes (Christin, 2017) to outright manipulations of the data sources that the automated systems rely on (Bachman, 2016; Wieczner, 2016).

However, our study differs from prior literature in important ways that contribute to this research area. A key novel insight the present work contributes is that occasionally inaccurate, partially “black-boxed” algorithmic systems can also contribute to trust and engagement if they reduce uncertainty. Our results show that when algorithms reduce critical uncertainties that workers face, they can contribute to the formation of trust and engagement among workers by enhancing perceptions of procedural justice. We found that the trust workers build in an algorithmic system is rooted in the perceived predictability and consistency of the system’s decisions. In addition, our findings show that strong collective positive emotions, such as those engendered by effective gamification, can further enhance workers’ engagement with the system and willingness to accept its occasional inaccuracies. This interpretation is consistent with findings in social psychology that positive moods are associated with increased prosocial behavior, a crucial indicator of engagement (Penner et al., 2005). We note that both enhanced perceptions of

procedural justice and positive collective emotions arose from workers' repeated interactions with each other and with the algorithmic system. It is unlikely that the effects on engagement would have been as strong in a setting where workers operated only in isolation from one another. Thus, our findings underscore the need for more in-depth studies of the ways in which collective sensemaking, interactional processes and enactment strategies shape workers' responses to technology in the workplace, a topic that has recently received increased attention from researchers in technology and organizations (e.g., Mazmanian, 2018; Beane and Orlikowski, 2015).

Network-based organizing

As network-based models of organizing proliferate and traditional organizations adopt elements of this form of organizing, understanding and developing theory around the factors that shape workers' experiences and relationships in network-based organizations is of rising relevance to organization scholars and practitioners alike. Using Phanish et al.'s (2014) criteria for assessing the novelty of an organizational form based on the novelty of its approaches to dividing up labor and integrating efforts, we present network-based organizations as a novel form of organizing that warrants further study from organizational scholars. In addition, our work suggests a potential extension of Phanish et al.'s (2014) criteria for organizational novelty to include the uncertainties that an organizational form presents workers with, as these uncertainties fundamentally shape workers' responses to management and control mechanisms. Paying greater attention to the uncertainties workers must navigate in novel organizational forms and the strategies they develop to do so is essential for the development of scholarly perspectives on modern organizing that balance managerial concerns with workers' lived experiences (e.g., Salehi et al., 2015; Kittur et al., 2013; Gray and Sunit, 2019).

Contributions to Group Engagement Research

The prior literature on group engagement has focused on the benefits of members' engagement for individual and organizational outcomes (e.g., Wollard and Shuck, 2011), but has paid relatively little attention to its antecedents. One notable exception is Tyler and Blader's (2003) group engagement model, which proposes that employees' perceptions of procedural justice with respect to interpersonal treatment and decision-making in the organization are critical for group engagement. However, the model's theoretical predictions have remained largely untested in empirical settings that involve algorithmically mediated relationships. We find that workers' perceptions of procedural justice, which are rooted in the algorithm's rule-based, consistent and predictable decisions and the sense of transparency these decisions added to the organization's otherwise opaque social landscape, play a critical role for the emergence of trust. However, our results also highlight the potentially important role of positive emotions in linking trust with engagement: We find that the strong positive shared emotions engendered by the algorithmic system's gamified aspects, together with workers' trust in the algorithmic system and the organization, supported high levels of engagement. While our data are not suitable for teasing apart causal relationships, future work should investigate the precise mechanisms underlying the relationships between positive emotions, algorithmic evaluations, trust, and engagement.

Uncertainty management

Uncertainty management theory deals with how employees cope with and respond to uncertainty in the workplace. This literature's core insight is that uncertainty is stressful, and that individuals' concerns around uncertainty are intimately linked to judgments of fairness (Lind and Van Den Bos, 2002; Van Den Bos and Lind, 2002). When people feel uncertain or focus on

uncertain aspects of their environment, they become especially concerned about the fairness of the treatment they experience (Lind and Van Den Bos, 2002; Van Den Bos and Lind, 2002). While organizational justice scholars have tended to argue that more fairness is always better (Colquitt et al., 2013), recent empirical work at the intersection of uncertainty management and organizational justice theory showed that people prefer being treated consistently unfairly over being treated sometimes fairly and sometimes unfairly, even if the average level of fairness experienced is greater in the latter than in the former scenario (Matta et al., 2017). In other words, within certain bounds, people seem to place greater importance on the degree of consistency in the fairness of treatment they can expect than on the average quality of treatment. Matta et al. (2017) conducted laboratory studies that showed that people who feel uncertain about the level of fairness they can expect to be treated with are more likely to experience exacerbated workplace stress, emotional exhaustion, and disengagement (Matta et al., 2017).

Our findings expand the link between justice variability and engagement these researchers uncovered. We find that even though freelancers did not generally think that the karma algorithm's decisions were particularly accurate or fair, they found these decisions to be highly consistent and predictable. Despite occasional discussions of and complaints about the inaccuracy of individual decisions the karma algorithm made, workers' responses to the karma algorithm for the most part focused around its dependable and unbiased properties, resulting in greater trust, engagement, and prosocial behavior. Consistent with findings by Matta et al. (2017), our observations suggest workers' preference for consistency in treatment over the average quality of treatment experienced.

Contributions to Ranking and Evaluation Research

Our findings also contribute to the literature on incentives, rankings, and evaluations. Behavioral economics studies on feedback and evaluations have found heterogeneous performance

effects (e.g., Kluger and Denisi, 1996). Rankings specifically seem to negatively affect employee effort and engagement, even if they are not tied to pecuniary incentives (e.g., Bandiera, Barankay, and Rasul, 2007; Barankay, 2010, 2011). For example, experiments showed that ranking employees based on their task performance elicited greater effort from those ranked at the very top but left the majority of employees feeling demotivated and disengaged (Gill et al., 2019). In addition, rankings systems where rewards are allocated based on rank have been found to crowd out employees' intrinsic motivation to engage in valuable discretionary behaviors that do not explicitly factor into the ranking, as employees focus on the codified behaviors that management has announced will receive recognition at the expense of other non-codified behaviors (see e.g., Gallus and Frey (2016) for a review). In part based on this research, many Fortune 500 companies, such as Microsoft in 2012, have shifted away from rankings systems to incentivize performance.

To our knowledge, the present study presents one of the first long-term scholarly investigations of workers' responses to and granular sensemaking processes around rankings in a real-world organizational setting. We find that rankings can have positive effects on employee engagement if they provide workers with more clarity on and transparency into the organization's social landscape, and if they help workers navigate uncertainties with respect to accessing jobs, forming relationships, and evaluating their own and others' performance. In our setting, workers also responded positively to the ranking system's gamification. In traditional organizations, where employees tend to face less uncertainty around accessing jobs, forming relationships and assessing their own and others' performance, rankings may not provide the same level of utility to workers, leading them to feel alienated rather than reassured by the rankings. Field experiments on rankings in other, less uncertain organizational contexts support this proposition: For example, in Barankay's (2011) field experiment on the performance effects of rankings in a furniture sales

organization, “all employees were well-informed about their current absolute performance and ability” even without the rankings, so the rankings did not have the uncertainty-reducing effects we observed in our study.

We propose that freelancers’ engagement levels and prosocial behaviors dropped when the karma system was discontinued because freelancers lost the karma system as a tool to navigate uncertainties. A potential alternative explanation for the decline in engagement comes from behavioral economics research, which has found that once individuals become used to receiving a reward for certain behaviors, their intrinsic motivation to engage in those behaviors dissipates (see e.g., Gallus and Frey (2016) for a review). As a result, they stop engaging in the behavior once the reward is no longer available. Because karma points functioned as a reward for certain prosocial behaviors, it is plausible that freelancers became accustomed to receiving karma awards in exchange for these behaviors, and that the incentive to engage in them diminished when karma was discontinued. Our data strongly suggest the importance of uncertainty reduction in the dynamics we observed, but we cannot exclude the crowding out effects based on our data. Future research can explore this question.

Boundary Conditions and Generalizability of Findings

Our case study represents an “extreme case” of a network-based organization, but we note that the uncertainties freelancers faced in this organization are becoming increasingly common in traditional organizations: In recent years, many organizations have begun to adapt remote working practices, have started to draw on hybrid workforces of full-time and contract employees, and increasingly rely on algorithmic systems to evaluate employees. As a result, the uncertainties that our study identifies as pertinent to workers’ experiences in network-based organizations are likely

to be relevant in workplaces in the future. At the same time, there are important bounds to the generalizability of our findings, which we outline below.

Quality of the algorithm's decisions and transparency into its policies. The extent to which algorithmic evaluation systems facilitate the formation of trust and engagement among workers is contingent on their perceived overall quality of the system's decisions. Occasional disagreements with the karma system's decisions notwithstanding, FlashCo freelancers thought that the system's decisions were acceptable in most cases. Moreover, we suspect that the public accessibility of the policies governing the karma system was key to its acceptance among freelancers. Had they not known what actions resulted in what types of karma decisions, they would likely have perceived the system as a "black box" and would have engaged in stronger resistance tactics akin to reports in the prior literature (Eubanks, 2018; Noble, 2018; O'Neil, 2016; Pasquale, 2015; Scholz, 2012; Zuboff, 2019; Danaher, 2016; Lebovitz et al., 2019).

Benchmark against which the algorithm's decisions are compared. In our setting, workers frequently compared the karma system to human managers who were either hypothetical or with whom they had worked in the past. Many freelancers preferred being evaluated by an algorithm that was "blind" to individual differences but made occasionally inaccurate decisions to being evaluated by a human whose average decision quality may have been higher, but who might have been biased against certain individuals in hard-to-predict ways. Future research should pay greater attention to the benchmarks people compare an algorithm's decisions against because their salient point of comparison shapes workers' responses in significant ways. These benchmarks are likely to vary as a function of workers' prior experiences interacting with human and non-human decision-making and evaluation systems at work and in private settings. For example, drivers on

ride-sharing platforms such as Uber may have never worked for human taxi-ride dispatchers and are thus unlikely to use them as points of comparison.

Workers' professional bias favoring algorithms. The network we studied consisted of software development professionals, about half of whom pursued active careers in technology companies, while the other half pursued software development freelance work full-time. Given their experience with and understanding of technological systems, it is probable that individuals in our sample were more trusting of algorithms, including the karma algorithm, than individuals from other, less technologically versed professions. Moreover, rankings and quantification are common in the software development community: Coding tournaments with real-time rankings (for example, kaggle.com) and software-specific knowledge-sharing platforms that rank contributors based on peers' upvotes (for example, stackoverflow.com) offer prominent examples. Because the freelancers in our sample were familiar with other online communities that relied on quantification and rankings, the algorithmic karma ranking system may have seemed less alienating to them than it would have to workers in other professions. Nonetheless, the karma system represented significant novelties even to the freelancers in our sample: All informants reported that karma was the first and only system they interacted with whose decisions were highly economically consequential and were based on algorithmic evaluations of their online behavior.

Workers' sense of agency. Compared with other gig work platforms, the workers in our sample exhibited relatively high levels of agency and self-confidence vis a vis the organization's headquarters. From their discussions in public Slack channels, it was evident that some freelancers held the perception that the organization needed them in order to turn a profit at least as much as they needed the organization as a source of work and income. In other words, many freelancers did not think that they were easily replaceable for the organization. This attitude provided the

foundation for freelancers' willingness to occasionally band together to push back against karma policies they did not agree with. Freelancers built a strong sense of solidarity for each other, giving rise to a sense of collective agency and "community ownership" of the network. This sense is less likely to emerge in settings such as Uber and Upwork, where workers do not interact with each other as they work. Workers' power, autonomy, and community-building may be important for engagement and prosociality around algorithmic evaluation systems.

CONCLUSION

The present study offers new theoretical insights into the processes underlying the formation of workers' engagement with their organization and peers when relationships are mediated by algorithms. Algorithms that perform managerial tasks are rising in prominence in network-based models of organizing, which manage large and distributed workforces at scale. We show that these novel work environments present workers with different uncertainties than traditional bureaucratic organizations, and that algorithmic management systems can enhance, rather than diminish, workers' engagement when they reduce critical uncertainties. Moreover, our findings highlight the importance of workers' collective sensemaking and emotional entrainment around these algorithmic systems in shaping prosocial behaviors. Our work opens up a framework for novel avenues of research in network-based models of organizing, including the emergence of community, identity, and prosocial behavior in algorithmically mediated contexts.

References

“Deloitte Pixel”

2020 Deloitte US. Retrieved May 27, 2020, from

<https://www2.deloitte.com/us/en/pages/operations/solutions/enterprise-crowdsourcing-solution-pixel.html>

“Artella”

n.d. Retrieved May 27, 2020, from <https://www.artella.com/>

“Smart websites at your service | B12”

n.d. Retrieved May 27, 2020, from <https://www.b12.io/>

“Factored AI Engineering”

n.d. Retrieved May 27, 2020, from <https://www.factored.ai/>

“FlashCo”

n.d. Retrieved May 27, 2020, from <https://FlashCo.com/>

“MVP Factory”

n.d. Retrieved May 27, 2020, from <https://mvpfactory.co/>

“Terminal.io - Where top engineers meet fast-growing companies.”

n.d. Retrieved May 27, 2020, from <https://terminal.io/>

“Topcoder | Design & Build High-Quality Software with On-Demand Talent”

n.d. Retrieved May 27, 2020, from <https://www.topcoder.com/>

“Toptal - Hire Freelance Talent from the Top 3%”

n.d. Retrieved May 27, 2020, from <https://www.toptal.com/>

“Turing - Build Faster Than Ever With Vetted Remote Developers”

n.d. Retrieved May 27, 2020, from <https://www.turing.com/>

“Distributed - Build software better and faster than ever before”

n.d. Retrieved May 27, 2020, from <https://distributed.co/>

“Comet - The best IT freelancers available now”

n.d. Retrieved May 27, 2020, from <https://www.comet.co/>

Aneesh, A.

2006 *Virtual Migration: The Programming of Globalization*. Duke University Press.

Bachman, R.

2019 “Want to cheat your Fitbit? Try a puppy or a power drill.” *Wall Street Journal*.

Retrieved from <https://www.wsj.com/articles/want-to-cheat-your-fitbit-try-using-a-puppy-or-a-power-drill-1465487106>

- Bahat, R.
2015, July 9 “Even The Org Chart Is Now Made Of Software | TechCrunch.” TechCrunch.
Retrieved May 27, 2020, from <https://techcrunch.com/2015/07/08/even-the-org-chart-is-now-made-of-software/>
- Bailey, D. E., F. Briscoe, A. Galper, K. C. Kellogg, P. Leonardi, and H. Lifshitz-Assaf
2019 “Is the Future of Work Already Here? The Changing Nature of Professionals’ Work and Professionalism.” *Academy of Management Proceedings*, 2019: 13597. Academy of Management.
- Bandiera, O., I. Barankay, and I. Rasul
2007 “Incentives for managers and inequality among workers: Evidence from a firm-level experiment.” *The Quarterly Journal of Economics*, 122: 729–773. MIT Press.
- Barankay, I.
2011 “Gender differences in productivity responses to performance rankings: Evidence from a randomized workplace experiment.” Unpublished Working Paper, The Wharton School.
- Barankay, I.
2010 “Rankings and social tournaments: Evidence from a field experiment.” University of Pennsylvania Mimeo, 15.
- Barrett, M., E. Oborn, and W. Orlikowski
2016 “Creating Value in Online Communities: The Sociomaterial Configuring of Strategy, Platform, and Stakeholder Engagement.” *Information Systems Research*, 27.
doi:10.1287/isre.2016.0648
- Beane, M., and W. J. Orlikowski
2015 “What Difference Does a Robot Make? The Material Enactment of Distributed Coordination.” *Organization Science*, 26: 1553–1573. INFORMS Inst.for Operations Res.and the Management Sciences.
- Bogost, I.
2015 “Why gamification is bullshit.” *The Gameful World: Approaches, Issues, Applications*, 65. MIT Press.
- Boudreau, K. J., and K. R. Lakhani
2013, April. “Using the Crowd as an Innovation Partner.” *Harvard Business Review*. Boston, MA. Retrieved from <https://hbr.org/2013/04/using-the-crowd-as-an-innovation-partner>
- Boyd, Danah, and K. Crawford
2012 “Critical questions for big data.” *Information, Communication & Society*, 15: 662–679. Routledge.
- Cameron, L.

- 2020 The Rise of Algorithmic Work: Implications for Organizational Control and Worker Autonomy. University of Michigan. Retrieved from <http://hdl.handle.net/2027.42/155277>
- Charmaz, K.
2008 “Reconstructing grounded theory.” *The SAGE Handbook of Social Research Methods*, 461–478. Sage Thousand Oaks, CA.
- Christin, A.
2017 “Algorithms in practice: Comparing web journalism and criminal justice.” *Big Data & Society*, 1–14.
- Christin, A.
2018 “Counting clicks: Quantification and variation in web journalism in the United States and France.” *American Journal of Sociology*, 123: 1382–1415.
- Colquitt, J. A., B. A. Scott, J. B. Rodell, D. M. Long, C. P. Zapata, D. E. Conlon, and M. J. Wesson
2013 “Justice at the millennium, a decade later: A meta-analytic test of social exchange and affect-based perspectives.” *Journal of Applied Psychology*, 98: 199. American Psychological Association.
- Colquitt, J. A., J. A. LePine, R. F. Piccolo, C. P. Zapata, and B. L. Rich
2012 “Explaining the justice-performance relationship: Trust as exchange deepener or trust as uncertainty reducer?” *Journal of Applied Psychology*, 97: 1–15.
- Danaher, J.
2016 “The threat of algocracy: Reality, resistance and accommodation.” *Philosophy & Technology*, 29: 245–268. Springer.
- Penner, L. A., J. F. Dovidio, J. A. Piliavin, and D. A. Schroeder
2005 “Prosocial Behavior: Multilevel Perspectives.” *Annual Review of Psychology*, 56: 14–15.
- Ederly, D., and E. Mollick
2008 *Changing the Game: How Video Games Are Transforming the Future of Business*. Ft Press.
- Edwards, R.
1979 “Contested terrain: The transformation of the workplace in the twentieth century.” Retrieved from <https://philpapers.org/rec/EDWCTT-2>
- Espeland, W. N., and M. L. Stevens
2008 “A sociology of quantification.” *European Journal of Sociology/Archives Européennes de Sociologie*, 49: 401–436. Cambridge University Press.
- Eubanks, V.

2018 *Automating Inequality : How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Publishing Group.

Faraj, S., S. Pachidi, and K. Sayegh
2018 "Working and organizing in the age of the learning algorithm." *Information and Organization*, 28: 62–70.

Fayard, A.-L., E. Gkeredakis, and N. Levina
2016 "Framing Innovation Opportunities While Staying Committed to an Organizational Epistemic Stance." *Information Systems Research*, 27: 302–323.

Fleming, J. H., and J. Asplund
2007 "Where employee engagement happens." *The Gallup Management Journal*, 3.

Gallus, J., and B. S. Frey
2016 "Awards: A strategic management perspective." *Strategic Management Journal*, 37: 1699–1714. Wiley Online Library.

Gill, D., Z. Kissová, J. Lee, and V. Prowse
2019 "First-place loving and last-place loathing: How rank in the distribution of performance affects effort provision." *Management Science*, 65: 494–507.

Glaser, B., and A. Strauss
1967 "The discovery of grounded theory." Weidenfield & Nicolson, London, 1–19.

Gray, M. L., and S. Suri
2019 *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books.

Herzberg, F.
1968 "One more time: How do you motivate employees." *Harvard Business Review Boston*, MA.

Herzberg, F.
1959 "The Motivation to Work." New York.

Horesh, R., K. R. Varshney, and J. Yi
2016 "Information retrieval, fusion, completion, and clustering for employee expertise estimation." 2016 IEEE International Conference on Big Data (Big Data): 1385–1393.

Horton, J. J.
2010 "Online labor markets." *International Workshop on Internet and Network Economics* Vol. 6484 LNCS: 515–522. Berlin, Heidelberg: Springer.

Jaros, S.
2010 "The core theory: Critiques, defences and advances." *Working Life: Renewing Labour Process Analysis*, 70–90. Palgrave Basingstoke.

- Kellogg, K. C., M. A. Valentine, and A. Christin
2020 “Algorithms at work: The new contested terrain of control.” *Academy of Management Annals*, 14: 366–410. Academy of Management.
- Kittur, A., J. V. Nickerson, M. S. Bernstein, E. M. Gerber, A. Shaw, J. Zimmerman, M. Lease, et al.
2013 “The future of crowd work.” *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*: 1301–1317.
- Kittur, A., L. Yu, T. Hope, J. Chan, H. Lifshitz-Assaf, K. Gilon, F. Ng, et al.
2019 “Scaling up analogical innovation with crowds and AI.” *PNAS*, 116: 1870–1877.
- Kluger, A. N., and A. DeNisi
1996 “The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory.” *Psychological Bulletin*, 119: 254. American Psychological Association.
- Lebovitz, S., N. Levina, and H. Lifshitz-Assaf
2019 “Doubting the Diagnosis: How Artificial Intelligence Increases Ambiguity during Professional Decision Making.” Available at SSRN 3480593.
- Lehdonvirta, V., O. Kässi, I. Hjorth, H. Barnard, and M. Graham
2019 “The Global Platform Economy: A New Offshoring Institution Enabling Emerging-Economy Microproviders.” *Journal of Management*, 45: 567–599.
- Leonardi, P., and N. Contractor
2018, November “Better People Analytics.” *Harvard Business Review*. Cambridge, MA. Retrieved from <https://empowerment.ee/wp-content/uploads/2018/11/Better-People-Analytics-Measure-Who-They-Know-Not-Just-Who-They-Are.pdf>
- Leventhal, G. S.
1980 “What should be done with equity theory?” *Social Exchange*: 27–55. Springer.
- Lind, E. A., and K. Van den Bos
2002 “When fairness works: Toward a general theory of uncertainty management.” *Research in Organizational Behavior*, 24: 181–224. JAI PRESS LTD.
- Lipsky, M.
2010 *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service*. Russell Sage Foundation.
- Matta, F. K., B. A. Scott, J. A. Colquitt, J. Koopman, and L. G. Passantino
2017 “Is consistently unfair better than sporadically fair? An investigation of justice variability and stress.” *Academy of Management Journal*, 60: 743–770.
- Mazmanian, M., and C. M. Beckman

- 2018 “‘Making’ your numbers: Engendering organizational control through a ritual of quantification.” *Organization Science*, 29: 357–379.
- McAfee, A., and E. Brynjolfsson
2017 *Machine, Platform, Crowd: Harnessing Our Digital Future*. WW Norton & Company.
- Miles, M. B., and A. M. Huberman
1984 “Drawing valid meaning from qualitative data: Toward a shared craft.” *Educational Researcher*, 13: 20–30. Sage Publications Sage CA: Thousand Oaks, CA.
- Noble, S. U.
2018 *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.
- O’Neil, C.
2016 *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Broadway Books.
- Orlikowski, W. J., and S. V. Scott
2014 “What happens when evaluation goes online? Exploring apparatuses of valuation in the travel sector.” *Organization Science*, 25: 868–891.
- Ott, B.
2007 “Investors take note: Engagement boosts earnings.” *The Gallup Management Journal*.
- Pasquale, F.
2015 *The Black Box Society*. Harvard University Press.
- Pine, K. H., C. Wolf, and M. Mazmanian
2016 “The Work of Reuse: Birth Certificate Data and Healthcare Accountability Measurements.” *iConference 2016 Proceedings*. doi:10.9776/16320
- Rahman, H. A.
2018 “From Iron Cages to Invisible Cages: Algorithmic Evaluations in Online Labor Markets.” *Work and Organization Studies Seminar*, MIT.
- Ramamurthy, K. N., M. Singh, M. Davis, J. A. Kevern, U. Klein, and M. Peran
2015 “Identifying Employees for Re-skilling Using an Analytics-Based Approach; Identifying Employees for Re-skilling Using an Analytics-Based Approach.”
doi:10.1109/ICDMW.2015.206
- Richman, A.
2006 “Everyone wants an engaged workforce how can you create it.” *Workspan*, 49: 36–39.
- Rosenblat, A.
2018 *Uberland: How Algorithms Are Rewriting the Rules of Work*. *Uberland: How Algorithms Are Rewriting the Rules of Work*. University of California Press.
- Salehi, N., L. C. Irani, M. S. Bernstein, A. Alkhatib, E. Ogbe, and K. Milland

2015 “We are dynamo: Overcoming stalling and friction in collective action for crowd workers.”
Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing
Systems: 1621–1630.

Scholz, T.
2012 *Digital Labor: The Internet as Playground and Factory*. Routledge.

Shestakofsky, B.
2017 “Working Algorithms: Software Automation and the Future of Work.” *Work and
Occupations*, 44: 376–423.

Spradley, J. P.
1979 *The Ethnographic Interview*. New York, NY: Holt Rinehart and Winston.

Strauss, A., and J. Corbin
1990 *Basics of Qualitative Research*. Sage publications.

Thibaut, J. W., and L. Walker
1975 *Procedural Justice: A Psychological Analysis*. L. Erlbaum Associates.

Thompson, P., and S. Vincent
2010 “Labour process theory and critical realism.” *Working Life: Renewing Labour Process
Analysis*, 47–69. Palgrave Macmillan Basingstoke.

Truelove, E.
2019 “Integrating the crowd into the firm production process: The critical role of guided
mobilization.” Working Paper. Massachusetts Institute of Technology.

Tyler, T. R., and E. A. Lind
1992 “A relational model of authority in groups.” *Advances in Experimental Social Psychology*
Vol. 25: 115–191. Elsevier.

Tyler, T. R., and S. L. Blader
2000 “Cooperation in groups: Procedural justice, social identity, and behavioral engagement.”
Essays in Social Psychology. New York, NY, US.

Tyler, T. R., and S. L. Blader
2003 “The Group Engagement Model: Procedural Justice, Social Identity, and Cooperative
Behavior.” *Personality and Social Psychology Review*, 7: 349–361.

Valentine, M., and R. Hinds
2019 *Algorithms and the org chart*. Working Paper. Stanford University.

Valentine, M. A.
2018 “Renegotiating spheres of obligation: The role of hierarchy in organizational learning.”
Administrative Science Quarterly, 63: 570–606.

Van den Bos, K., and E. A. Lind

2002 “Uncertainty management by means of fairness judgments.” Academic Press.

Varshney, K. R., V. Chenthamarakshan, S. W. Fancher, J. Wang, D. Fang, and A. M. Mojsilović

2014 “Predicting Employee Expertise for Talent Management in the Enterprise.”
doi:10.1145/2623330.2623337

Walz, S. P., and S. Deterding

2015 *The Gameful World: Approaches, Issues, Applications*. MIT Press.

Wieczner, J.

2016 “Fitbit Users Are Finding Creative Ways to Cheat.” *Forbes*.

Wollard, K. K., and B. Shuck

2011 “Antecedents to employee engagement: A structured review of the literature.” *Advances in Developing Human Resources*, 13: 429–446.

Zuboff, S.

2019 *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Profile Books.

FIGURES

Figure 1: Number of contractors who sent messages (Panel A) and total number of messages sent (Panel B) in a network-wide public channel per month

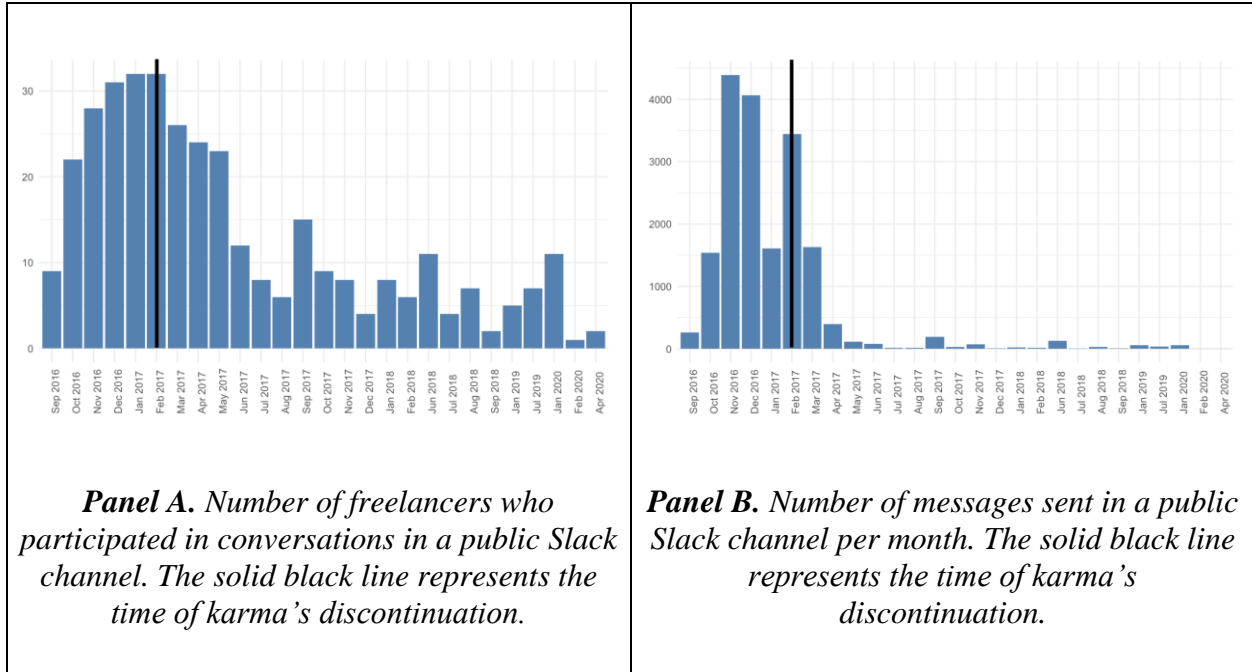
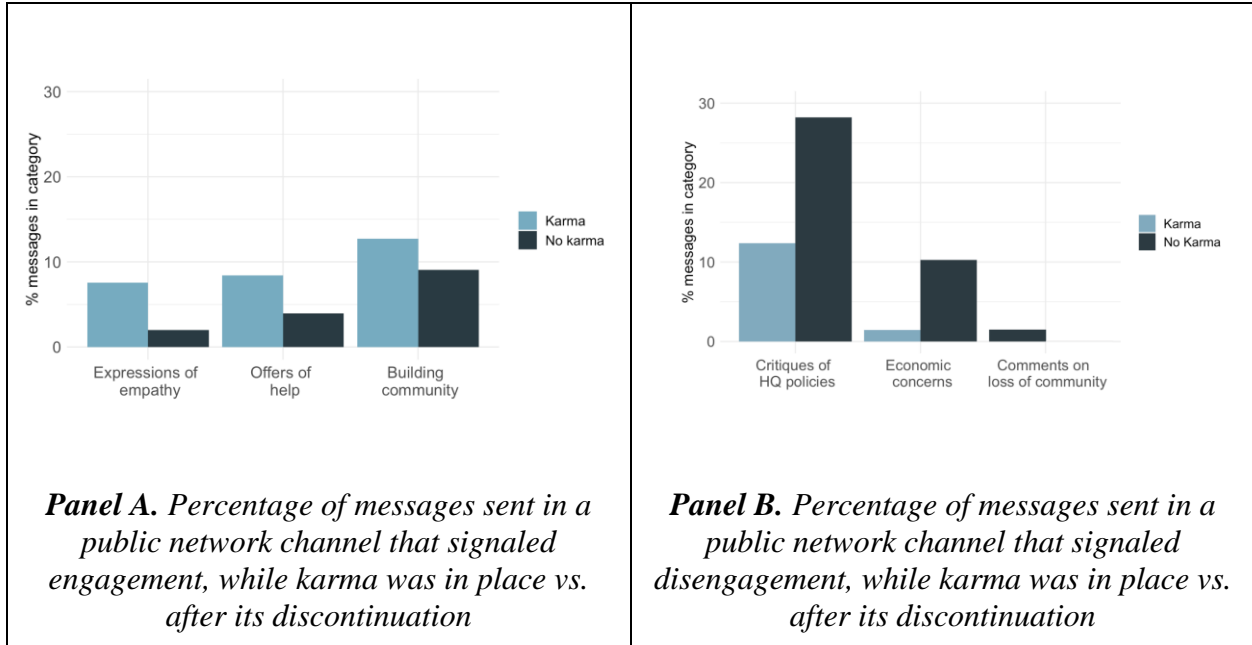


Figure 2: Comparison of substantive content of Slack messages exchanged in a public network channel while the karma system was in place versus after it was discontinued



Note: Y-axes represent the percentage of messages that were qualitatively classified as belonging to each primary category, calculated out of the total messages sent in the periods during which karma was in place and after it was discontinued, respectively. Each message was coded as belonging to only one primary category. Panels A and B show message categories that signaled engagement and disengagement, respectively.

TABLES

Table 1: Sources of uncertainty freelancers confronted in the network-based organization

	Illustrative data
Starting up	<p>Freelance project manager 4, in an interview: <i>“This isn’t [DIFFERENT ONLINE LABOR PLATFORM]. You’re getting paid flat rate, and you kind of need to understand how FlashCo works rather than just coming in with your own ideas about how freelance works.”</i></p> <p>Freelance project manager 7, in a public Slack channel: <i>“PMing on FlashCo is not like PMing at other places. Nobody tells you that... That isn’t apparent at first”</i></p> <p>Headquarters staff member 1, in an interview: <i>“Freelancing generally is a very hustle culture, and FlashCo is special even in that world. I do think people [who succeed on FlashCo] are pretty savvy ... about what they have to do in order to get certain outcomes. People who are a bit slower moving ... they start learning that it’s an uncertain world”</i></p>
Accessing jobs	<p>Freelance project manager 6: <i>“Getting the first job wasn’t that hard but getting a good one could be tricky. Good jobs aren’t served to you on a silver platter. As a freelancer on FlashCo, you’ve got to figure out the little mechanics for how to get what you want.”</i></p>
Forming relationships	<p>Freelance project manager 8, in a public Slack channel: <i>“As a [project manager], sometimes we just need to be able to trust who’s joining the team. That’s why we build relationships with a few [freelancers] [...] It’s hard to trust people in the network sometimes, man...”</i></p> <p>Freelance developer 5, in an interview: <i>“It’s all about knowing the right people and teaming up with them. If someone on your team abandons the project or leaves you hanging, that’s really bad. So, you’ve got to get to know the right people, know who to trust, and make friends with them.”</i></p>
Gauging own and others’ status and performance	<p>Freelance developer 10, in an interview: <i>“In freelance work, knowing your relative market value at all times is really important. If you think you’re better than you are, you ask for too much, and you don’t get jobs. If you undersell yourself, you leave money on the table and people see that. Either way, it doesn’t look good.”</i></p> <p>Freelance developer 5, in an interview: <i>“You always have to keep tabs on how good you are because this industry moves really quickly. If you</i></p>

don't do that, you don't know where your learning gaps are, and then you can't keep up with the industry."

Freelance project manager 9, in an interview: *"I was always careful to avoid the mistake of working with [team members] who aren't 100% professional and committed. Normally, you can ask around about someone's reputation and get a sense for that... But in a huge network with a thousand people, that is hard. You really need some other indicator of quality to help solve that problem"*

Table 2: Freelancers used Karma to make sense of the uncertain landscape

	Illustrative data
Accessing jobs	<p>Freelance developer 1, in an interview: <i>“It was very clear [how to get good work]: you get karma, you get [projects]. [...] there was a clear path forward for you.”</i></p> <p>Freelance developer 5, in an interview: <i>“Where karma was most useful for me was whenever I was looking for new [projects], ... you had to kind of develop a relationship with a sales manager. You can point to, ‘Hey, here’s my karma score.’ And that would be a quick qualifier that they [project managers] will look at ... to verify that you were actually worth their time or would be a good fit for different [projects] they had coming in the pipeline.”</i></p>
Forming relationships	<p><i>Bonding over karma</i></p> <p>Freelance developer 1, in an interview: <i>“Back then [when karma was around], you made a friend because you had something in common: You both wanted karma, you both wanted to work on cool projects. You would help each other out.”</i></p> <p>Freelance developer 11, in an interview, on karma’s role as a versatile and safe topic of conversation: <i>Karma meant “there was always another topic to talk about ... It was a nice tool to get to know your peers given that everyone on FlashCo is remote”</i></p> <p>Freelance developer 1, in an interview: <i>“[Karma] wasn’t a perfect system ... But there was a relational model there that came in conjunction with it, where once you [use karma to] make friends, you can get [projects].”</i></p> <p>Humorous karma talk spawned casual interactions, an important source of bonding. From a public Slack channel: Freelance developer 13: <i>“[Freelance developer 14], the grocery store guy insists on not accepting karma for my food, unacceptable”</i> Freelance developer 14: <i>“I’d love to see a video of you going around to different establishments and being like, ‘Do you take karma?’”</i></p> <p><i>Learning and shaping norms</i></p> <p>Freelance project manager 5, in an interview, recounted receiving and using karma to reward acts of helping that were not required by FlashCo’s formal rules: <i>“I remember being on the receiving end and also giving karma for when somebody did a really good job, helped me with troubleshooting an issue or something. I would award them a couple of karma points. And I think that was definitely the norm. [...] I thought it</i></p>

	<p><i>was kind of cool. [Giving or receiving karma] made me feel good about the day.”</i></p> <p>Freelance project manager 5, in an interview: <i>“Most people tended to do good things. I mean folks who did bad things had karma removed, but I think that was a good sort of... I think they saw the value of the incentive system.”</i></p> <p><i>Finding mentors</i></p> <p>Freelance developer 5, in an interview: <i>“When I came in, I looked at people who had really high karma as a reference that they're the ones who know how to get work reliably and are making a living through FlashCo. So, I wanted to be able to connect with them and learn from them ... [karma] was useful for me to distinguish who would be a good person to connect with as a mentor.”</i></p>
<p>Gauge own and team’s performance</p>	<p>Freelance project manager 4, in an interview: <i>“[Karma was] great for people to be able to establish a rank amongst themselves, where they can still be members of the community but they can also find who to learn from and place people socially within FlashCo.”</i></p> <p>Freelance developer 2, in an interview: <i>“I think the main reason karma was good was just that there's some kind of ranking system that is based on a key variable that is public [rather than] secret.”</i></p> <p>Freelance developer 6, in an interview: <i>“People respected you if you had a good karma count [...] it was the reputation, even if they didn't know anything else about you, especially at 1000 plus karma. People were really proud of that because [a karma score that high means] you've been around for a while.”</i></p>

Table 3: Karma supported freelancers’ perceptions of procedural justice

	Illustrative data
Karma increased perceived quality of decision-making	<p>Freelance developer 1, in an interview: <i>“With karma, I knew exactly what I had to do to get where I needed to be to get the jobs I want.”</i></p> <p>Freelance developer 2, in an interview: <i>“I always felt like I knew why certain decisions were made [with respect to staffing]. It was all based on karma. I had the sense that I knew what was going on.”</i></p> <p>Headquarter staff member 1, in an interview: <i>“[Freelancers] were savvy about how to get the outcomes they wanted ... Karma made it clear what they needed to do to succeed. They appreciated that.”</i></p>
Karma increased perceived quality of interpersonal treatment, including prosocial behavior	<p>Freelance project manager 5, in an interview: <i>“Karma was a nice way to bond with each other. It made people do good things. ... [Team-level karma decisions created] a nice dynamic because the whole team had to work together.”</i></p> <p>Freelance developer 8, in an interview: <i>“I actually used it [karma] quite a bit [on my teams]. I used the positive [karma] ... a lot. It did generate kind of a positive feedback loop. We just had a good time working together.”</i></p> <p>Freelance developer 1, in an interview: <i>Karma led to “more conversation, people helping each other... Karma was definitely an incentive factor to help each other get past obstacles or being stuck.”</i></p> <p>In public Slack channel, while karma was in place, freelance developer 15 described the network’s culture: <i>“We have a great culture within the community. We help each other out. It’s a safe place to ask questions and admit you fucked up. We [have] a lot of mutual trust. We are all very open with each other. We venture outside of the bubbles we work in.”</i></p> <p>In a public channel, freelance developer 15 explained how he used karma to express appreciation for others: <i>“I try to always give people a little ... karma if they help with something. I know it may not be much, but it shows I value their time.”</i></p>

Table 4: Karma produced strong collective emotions

	Illustrative data
Gamification elicited strong collective emotions	<p>Freelance developer 8, in an interview: <i>“I’m very, very competitive. I very much wanted to earn karma, it actually got to a point where earning karma was really more important to me than earning money.”</i></p> <p>Freelance project manager 5, in an interview: <i>“I... liked that you got all of this reinforcement when you're working [...] it feels like a game, and you're constantly rewarded as you do good things”</i></p> <p>Freelance designer 3, in a public Slack channel: <i>“It’s nice to get the reward when you are on a project... I love that feel... karma has such a good vibe”</i></p> <p>Freelance developer 15, in a public Slack channel: <i>“Is it bad that I know karma doesn't mean anything, but it still makes me feel all warm and fuzzy when mine goes up!?”</i></p> <p>From a team’s Slack channel, after the team received a karma award: Freelance project manager 11: <i>“All the yapping got us sprinkled [with karma]. My goal is to get so much micro karma that I get back to even numbers. Then I stop [happy emoji]”</i></p> <p>From a public Slack channel: Humorous conversation around karma facilitates social bonding between freelancers: Freelance developer 14: <i>“You BASTARD! You bumped ahead of me in the great karma war.</i> Freelance developer 16: <i>“Early :birds catch the fake karma worms”</i> Freelance developer 14: <i>BTW did you hear? [Freelance project manager 10] is karma queen again.</i></p> <p>From a project team’s Slack channel, after the team received a karma award: Freelance developer 17: <i>“We 're going to Vegas!! [happy emoji]”</i> Freelance project manager 11: <i>“This could be a life changer.”</i></p> <p>From a project team’s Slack conversation: After the project manager has worked through a tough dispute with the client, the team earns back some karma points it had lost. The freelance developer addressed the PM: <i>“[I’m] in the middle of a run. ... I had to stop and say thank you for all that you have done. These karma points are so significant to us.”</i></p>
Solidarity through collective action	<p>Freelance project manager 6, in an interview: <i>“[The network company’s founder] would send an email saying, ‘This [new karma policy] is</i></p>

getting rolled out.’ Sometimes there was no notice, and that was even worse. But he learned to at least try to announce it, and then we’d all go, ‘Hmm... I don’t know about this!’, and then we’d talk about it with each other, and then it would get rolled out and then we would explode, right.”

Freelance developer 8, in an interview: *“I worked hard to do well in the karma system and when I thought that the game wasn’t working the way I thought it should, I was verbal about it.”*

Freelance developer 2, in an interview: *“I remember a lot of people being really annoyed by karma in Slack channels and arguing about how fair it was, saying that they got a score change, or some bot took points from them and that it wasn’t fair. But, again, the people that were there before, they had amassed some decent karma, they loved it.”*

Freelance developer 15 to other freelancers in the public Slack channel, as they prepare to submit a list of requests for karma policy changes to headquarters: *“If we hit hard, we can illicit change. But we have to be focused. At the end of the day, we control the money.”*

Table 5: Decline in engagement following discontinuation of karma

	Illustrative data
Decline in participation rates	<p>Freelance developer 1, in an interview, on the changes following karma’s discontinuation: <i>“There used to be so much more interaction on Slack. Conversation, people helping each other, people giving jobs to each other, asking, ‘Can you do this for me?’ It used to be a lot more than it is [now].”</i></p> <p>Freelance project manager 1, in an interview: <i>“[When karma was in place], the community was a lot more boisterous, it was a lot more active and chatty ... I would definitely say that karma brought about a heightened sense of engagement on the platform. It was a reference point. People were talking about karma.”</i></p>
Decline in sense of community	<p>Freelance project manager 4, in an interview: <i>“I think there's a bit of a community aspect that is lost [without karma]... It's a pretty big loss. I'd say it's significant.”</i></p> <p>Freelance project manager 1, in an interview: <i>“In terms of the camaraderie ... the platform is nothing like [what it was when karma was in place] now... I think the general vibe on the platform is more muted. People are more focused on just delivering their [projects].”</i></p> <p>Freelance developer 1, in an interview: <i>“Yes [the community on FlashCo was something I greatly enjoyed when karma was around]. I was part of [a public community Slack channel]... I talked to them every day... I had friends. And now I don't know where they are anymore.”</i></p> <p>From a public Slack channel: Freelancers commented on the loss of engagement following karma’s removal: Freelance developer 15: <i>“[This public channel] has turned into a ghost town. [sad emoji] No messages in [in this channel] for over a week...”</i> Freelance developer 18: <i>“It's sad man [disappointed emoji]”</i></p> <p>Freelance developer 4, in an interview: <i>“There are many people that are not on FlashCo anymore [as a direct result of the loss of karma]”</i></p>
Uncertainties resurfaced	<p>Freelance developer 1, in an interview: <i>“Karma allowed me to hustle my way to where I needed to get to and that's why I felt efficient on FlashCo, whereas now [that karma is gone], I don't understand the hustle.”</i></p>

<p>Lower perceived transparency and lower levels of trust</p>	<p>In aftermath of a town hall meeting around the time withdrawal of karma was announced, freelance developer 15 said in a public channel: <i>“Trust is incredibly important to me in my work. I have earned their [headquarters’] trust, but they have yet to earn mine. [Headquarters] willfully took it, threw it in the mud, and jumped on it with cleats ... And then I’m supposed to listen when they say ‘trust us.’”</i></p>
<p>Lower perceived quality of procedural justice with respect to interpersonal treatment</p>	<p>When karma is discontinued, freelancers become upset at the quality of treatment they experience from headquarters. From a public Slack channel: Freelance developer 19: <i>“They [headquarters] just need to figure out a better system and treat the freelancers the way we deserve to be treated...”</i> Freelance developer 20: <i>“It’s like they really don’t trust their own engineers for anything. God... this mentality they have permeates everything they do.”</i></p>
<p>Lower perceived quality of procedural justice with respect to decision-making</p>	<p>After karma was discontinued, freelancers demanded more transparency from headquarters into how decisions were made but did not feel that their demand was met. From a public Slack channel: Freelance developer 21: <i>“Definitely more transparency is needed [...] but I feel like HQ won’t share that”</i> Freelance designer 3: <i>“Not sure why they [headquarters] keep avoiding the transparency issue. Look, they have no obligation to be transparent if they don’t want to. So, if that’s their stance, then own it. It’s the wishy-washy crap that is sinking HQ.”</i> Freelance developer 15: <i>“I have been one of the strongest proponents of FlashCo and today they lost my support [confused emoji]”</i></p>