Ink: Increasing Worker Agency to Reduce Friction in Hiring Crowd Workers

NILOUFAR SALEHI and MICHAEL S. BERNSTEIN, Stanford University

The web affords connections by which end-users can receive paid, expert help—such as programming, design, and writing—to reach their goals. While a number of online marketplaces have emerged to facilitate such connections, most end-users do not approach a market to hire an expert when faced with a challenge. To reduce friction in hiring from peer-to-peer expert crowd work markets, we propose Ink, a system that crowd workers can use to showcase their services by embedding tasks inside web tutorials—a common destination for users with information needs. Workers have agency to define and manage tasks, through which users can request their help to review or execute each step of the tutorial, for example, to give feedback on a paper outline, perform a statistical analysis, or host a practice programming interview. In a public deployment, over 25,000 pageviews led 168 tutorial readers to pay crowd workers for their services, most of whom had not previously hired from crowdsourcing marketplaces. A field experiment showed that users were more likely to hire crowd experts when the task was embedded inside the tutorial rather than when they were redirected to the same worker's Upwork profile to hire them. Qualitative analysis of interviews showed that Ink framed hiring expert crowd workers within users' well-established information seeking habits and gave workers more control over their work.

CCS Concepts: • **Human-centered computing** → *Computer supported cooperative work*;

Additional Key Words and Phrases: Social computing, crowdsourcing

ACM Reference format:
Niloufar Salehi and Michael S. Bernstein. 2018. Ink: Increasing Worker Agency to Reduce Friction in Hiring Crowd Workers. *ACM Trans. Comput.-Hum. Interact.* 25, 2, Article 10 (April 2018), 17 pages.
https://doi.org/10.1145/3177882

1 INTRODUCTION

Despite years of research and industry effort, peer-to-peer crowd work marketplaces are still not common destinations for end users. While 22% of Americans have contributed to an online fundraising project, and 15% have used ride-hailing apps such as Lyft, only 4% have hired someone online to do a task or errand[1] (e.g., from Amazon Mechanical Turk, Upwork, or TaskRabbit). For example, writing and editing are common but difficult information work activities, and editors are readily available on expert crowd markets such as Upwork. However, few people hire help

[1] http://www.pewinternet.org/2016/05/19/the-new-digital-economy/.

This work is supported by the National Science Foundation, under Grant IIS-1351131.
Authors' addresses: N. Salehi and M. S. Bernstein, Gates Computer Science, 353 Serra Mall, Stanford, CA 94305; emails: {niloufar, msb}@cs.stanford.edu.

https://doi.org/10.1145/3177882

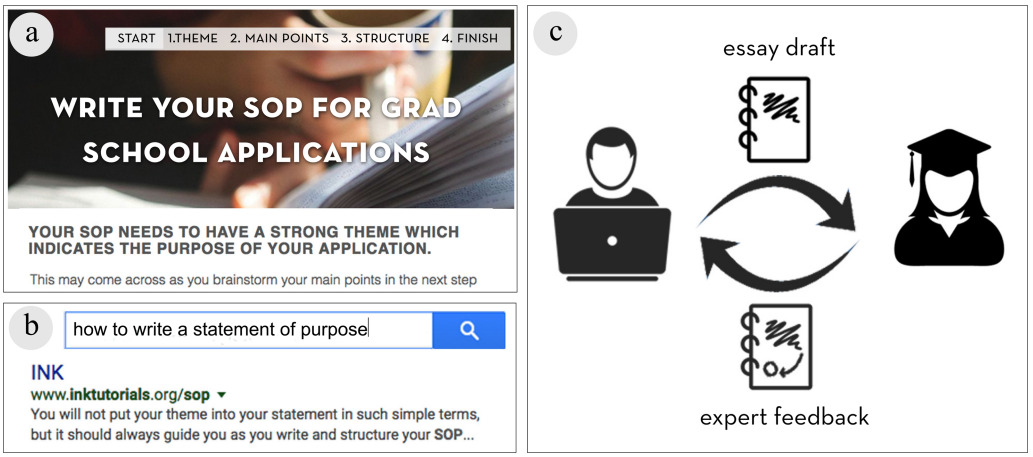

Fig. 1. (a) Crowd experts create Ink tutorials and augment them with tasks. (b) Users find the tutorial and tasks through web search. (c) Crowd experts provide paid help.

online with writing. Crowd experts are similarly available but vastly underutilized across domains, such as programming, design, and legal services.

Why is there such a difference between the availability of crowd expertise and end-user utilization of that expertise? This article pursues a thesis that end users are not always aware of available expertise, and even when they are aware, the burden of crafting an effective job post and searching for the right people to hire can be prohibitive. We refer to this accumulation of opposing forces as end-user *friction* in expert crowd work. We suggest that this up-front friction of searching, authoring, and hiring will together trigger prompt loss aversion (Tversky and Kahneman 1992), where the immediate and certain time cost now outweighs the potential benefit later.

We propose that crowd work hiring surfaces can be designed to reduce friction by giving workers more agency to manage work and signal how their help may fit into end users' needs. We draw on the insight that end users do not go to marketplaces when they seek help, instead they search for help on search engines (Brandt et al. 2009). Tutorials in particular are highly popular on search result pages, and are familiar to end users (Lafreniere et al. 2013). We thus introduce a system, *Ink* (Figure 1), where crowd experts author web tutorials for complex goals and showcase their services as embedded tasks inside those tutorials. Crowd experts use Ink to create tutorials that guide the reader through the steps of achieving their goal (e.g., writing a Statement of Purpose (SoP) for graduate school) while offering to perform paid tasks such as giving feedback on work or completing a step of the task for the user. Ink tutorials look like traditional how-to guides on the web, making them familiar to web users, easy to share, and attractive to search engines. However, unlike traditional web tutorials, Ink tutorials contain hooks for hiring the author to assist with some steps. As a result, users just need to upload their documents and submit a task request. Ink tutorials cover diverse tasks such as making a compelling personal homepage or practicing for an interview.

We executed a field deployment of Ink publicly on the web for one year. We initially hired three crowd experts to author and host a seed set of tutorials on the site. Through this deployment, 168 people hired Ink's crowd experts for 190 tasks of $20.30 (s.d. = $16.70) on average. Fully 80% of these users had never hired anyone online prior to using Ink, suggesting that Ink tutorials reach a different category of users than traditional crowdsourcing marketplaces. We sent a short survey

to all ink users and conducted longer interviews with 16 users who had hired from the tutorials in one month. We also conducted semi-structured interviews with the experts who had created the tutorials. Users reported that they had a specific, context dependent goal that brought them to Ink and most had not hired anyone online for a task prior to Ink. Crowd experts reported that they felt more power to set expectations as managers of work rather than applicants to work posted by others.

Finally, we performed a field experiment to isolate the effect of embedding the pre-defined crowd tasks inside the tutorial rather than directing visitors to a marketplace where they could hire help. We randomly assigned half of Ink visitors to a condition where instead of the embedded Ink task, they saw a link to the same worker's Upwork profile. The Ink condition resulted in 32 task requests, while the Upwork condition resulted in only 1 request ($\chi^2(1, N = 33) = 29.1, p < 0.05$).

This article identifies causes of friction in hiring help from peer-to-peer expert crowd work markets and introduces a platform that resolves these barriers by giving workers the tools to present their services in a pre-defined and familiar format within users' regular information seeking paths. Toward this end, we present a system, a longitudinal deployment, and a field experiment. We view Ink as a step toward a future where expert crowdsourcing more fully integrates into end users' lives as a means for receiving services, mentorship, and expert feedback.

2 RELATED WORK

We build on prior work that engages a crowd workforce to complete tasks on demand (Kittur et al. 2013; Retelny et al. 2014; Suzuki et al. 2016). Specifically, we study peer-to-peer crowd markets where users can purchase services from crowd workers, often experts (Retelny et al. 2014).

2.1 Crowd Work

A large body of HCI research has focused on crowd-powered systems, articulating software visions that make crowd intelligence available to end users through domain-specific applications such as word processors (Bernstein et al. 2010) and visual question answering (Bigham et al. 2010). In the domain of peer-to-peer crowd work markets researchers have studied the impacts of crowdsourcing on the workers and organizations that employ them (Gray et al. 2016; Irani and Silberman 2013; Martin et al. 2014). However, there has been less focus on the design of platforms for hiring online work and on the challenges that regular end users face when attempting to hire from crowd work markets.

Researchers have proposed that the design of alternative crowd work platforms can be an opportunity to change the face of crowd work (Kittur et al. 2013), in particular the unbalanced power relations between workers and requesters (Irani and Silberman 2013; Kittur et al. 2013). More recent work has attempted to redesign crowdwork platforms for various objectives. For instance, worker led reputation systems have proven to outperform traditional platform based systems (Whiting et al. 2017). We build on this research to examine the possibility of designing crowd work platforms that reduce friction for the user by giving the worker more agency.

Online workers are legally considered independent contractors, the implication being that these workers have full control over where, when, and how much they work (Lee et al. 2015). However, the systems that manage their work often direct workers' collective power into computational infrastructure where their roles and abilities are specific and limited (Irani and Silberman 2013). As a result, some online workers feel that they have lost agency in managing their jobs (Lee et al. 2015). With Ink, we seek a more flexible, open-ended role—workers create their tasks and define the circumstances of their employment. We suggest that increased authority and flexibility will enable crowd workers to enhance their online careers.

Table 1. The Design Space of Hiring Surfaces for Crowd Work

Task creation	Created by user (e.g., AMT)	Pre-defined by worker (e.g., Fiverr)	Pre-defined in system (e.g., Soylent)
Task situation	Generic hiring platform (e.g., Upwork, Fiverr, TaskRabbit)	Tasks embedded within the context of the user's workflow	User unaware of crowdsourced backend (e.g., Facebook spam filter)
Task type	Well defined, repeatable, and common (e.g., Fiverr)	Unclear goals that require negotiation (e.g., Upwork)	Niche domains (e.g., Upwork, AMT)

Note: Ink's choices are highlighted.

2.2 Tutorials and Finding Help Online

Tutorials are a major source of structured help online. They are prevalent on the web today, covering a variety of skills and domains. Web-based tutorials serve the following main roles for their readers: in-task help for immediate, well-defined needs; a means to proactively expand a user's skills; and an opportunity to shadow experts (Lafreniere et al. 2013). Ink tutorials maintain these benefits, but open the opportunity for a short-term tutoring relationship. Users of web tutorials frequently need further help, often resulting in comment sections that evolve into support forums (Bunt et al. 2014; Lafreniere et al. 2013). Ink allows users to hire tutorial authors for personal help.

Q&A websites such as Stack Overflow are a second option for connecting people who need help to experts who can help them (Adamic et al. 2008). Volunteers on these platforms favor responding to questions that will benefit the community as a whole (Mamykina et al. 2011) and often react negatively to repeat questions. While valuable, this style of Q&A leaves out users who need individual, non-generalizable help such as feedback on code style. With Ink, we focus on these individual, non-generalizable use cases that need support. Complementing these volunteer systems with paid crowd services enables users to connect to an expert when they need personal help. Paid Q&A systems such as Google Answers have also offered personal information seeking services for users (Chen et al. 2010). However, these services face the same discovery barrier as generic crowd work platforms.

3 INK

Ink reduces friction for users by creating the conditions for expert crowd workers to frame their services in a context familiar to users. Expert crowd workers can use Ink to create tutorials and thread in their services as pre-defined tasks. Table 1 summarizes the design space of hiring surfaces for crowd work platforms. Our choices in designing Ink are highlighted. We have made the following design decisions in line with our goal of reducing friction in adopting expert crowd work by regular web users:

Task creation. Tasks may be created by the user or they may be pre-defined by either a crowd worker or by system designers of a crowd-powered system. When users create tasks, they benefit from high flexibility to describe their individual needs and to find a worker with the right skills to perform the task. For example, systems such as Upwork, AMT, and TaskRabbit leave it up to the user to create and submit any task, starting from a blank form. However, too much flexibility can create choice paralysis and make the whole task too difficult to start (Iyengar and Lepper 2000; Weeks 2004). On the other hand, when tasks are pre-defined they lift the burden of creating the task from users, making the services easier to adopt. To cover a wider range of tasks, systems can enable

crowd workers to create tasks. Expert workers can identify opportunities and create the tasks that they would like to work on. Fiverr is an example of such a system that hosts a large database of tasks that workers have created. With Ink, we argue that workers have more information about how their services may be of use to users; therefore, we enable workers to create tasks. We will argue that creating and managing tasks gives workers more agency in the relationship.

Task situation. A crowdsourcing task may be posted on a generic hiring platform, embedded within the context of a user's workflow, or remain in the backend unknown to the user. Common crowd work platforms such as Upwork, AMT, and Fiverr are generic marketplaces for tasks. To utilize these tasks, users need to recall an opportunity, exit their current workflow, and approach a generic market to meet their needs. This process relies on users both having the information that their need may be met on an online labor market and recalling it at the right time. Ink instead turns the cognitive task of the user from recall to recognition. To this end, we have designed Ink based on tasks that are embedded within the context of the user's workflow. This design decision ensures that users will come across Ink tasks within their regular path and will recognize the opportunity to utilize those tasks.

Task type. Tasks may be well-defined and repeatable, have unclear goals, or be in niche domains. Tasks that do not have clear goals require negotiation and back-and-forth communication between the user and crowd worker (Salehi et al. 2017). Platforms such as Upwork support this type of communication. Tasks may also be in niche domains, in which case the user will need to seek workers who fit the specific needs of the task. Generic crowd marketplaces such as AMT and TaskRabbit make this possible. A third class is tasks that are well defined, common, and repeatable. These tasks do not require authoring every time they are instantiated, instead they can be predefined and reused. Their popularity also makes it worthwhile for workers to endure some upfront costs to make these tasks accessible to a larger population. This will significantly reduce the barrier to their adoption; therefore, Ink supports this type of task.

4 SYSTEM

Ink supports expert crowd workers in creating step-by-step tutorials on the web. In doing so, it encourages the worker to insert their own services in the tutorial, for example, to give feedback on drafts or to complete a step of the tutorial for the user. Ink tutorials are grounded in a series of steps (Figure 2(a)). For example, a website design tutorial for personal homepages includes the following: (1) gather and organize content, (2) create a wireframe, (3) add visual design, (4) code HTML and CSS, and (5) host the page.

Ink tutorials offer tasks on some steps, where the experts offer paid individual help (Figure 2(b)). For example, for the tutorial in Figure 2(a), step (4) offers—"Upload your content and visual design and I will code your website in HTML and CSS" for $30. Or step (5)—"Ask me any questions you have about hosting", offered for $1. These task interfaces specify the logistics, any required inputs, and the price. Unlike most crowdsourcing platforms, where the user specifies the task description and requirements, here the workers write and frame the task that they will complete for the user within the context of the user's workflow. Ink inserts each task after its relevant step and also displays all of the tasks at the top of the tutorial, making them more salient to users with targeted needs.

To make a request to the worker, a user fills out the required information and submits the form. Ink sends an email to both the user and the expert with instructions on how to proceed. These instructions vary by task and include payment instructions (e.g., PayPal), time estimates for receiving results, and scheduling information for online coaching sessions. Users are instructed to pay before the service is rendered in order to protect experts. Ink provides a template for this email

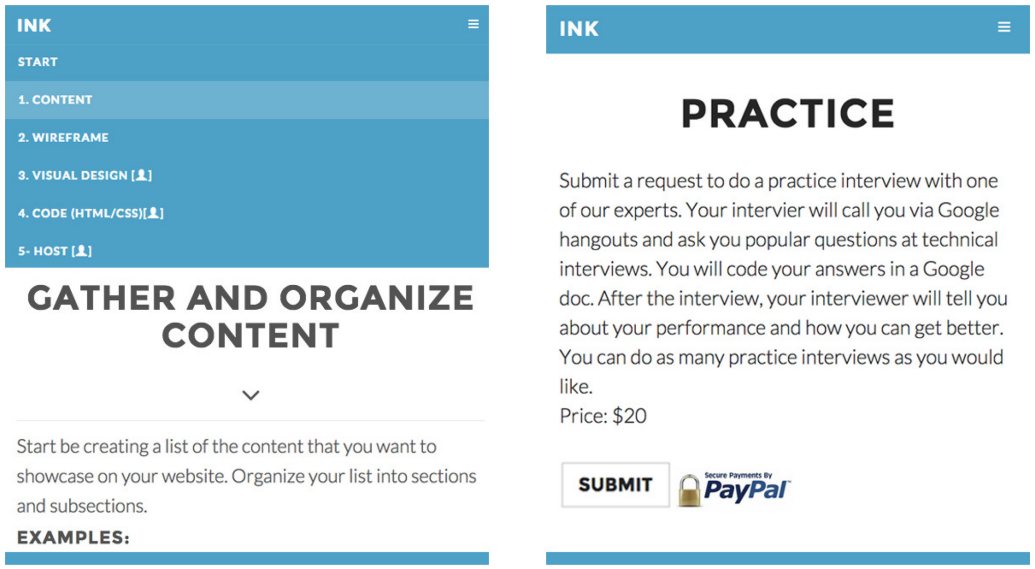

Fig. 2. (a) Expert crowd workers create Ink tutorials to guide users step-by-step to their final goal. This tutorial, by a web designer on Upwork, helps create a personal website. (b) Experts equip Ink tutorials with crowdsourced tasks that offer to help users, for example, by practicing a programming interview.

and the expert who created the tutorial can customize the message to meet their needs. Once the task is done, the user can return to Ink to leave a rating that is visible to future visitors. Ink displays ratings for each tutorial to signal quality and trustworthiness.

Ink supports online experts in creating these crowd-augmented tutorials. To do so, it provides templates for tutorials and tasks, hosts the tutorials, administers a reputation system, and distributes tasks to workers.

4.1 Tutorial Templates Guide Experts' Writing

Crafting an effective message and storefront is not easy. To help, Ink hosts templates for creating tutorials and task descriptions. The tutorial templates consist of a variable number of steps toward a final goal. Each step consists of a title, content, and optional images and examples. Ink displays these tutorials within a Twitter Bootstrap template that is consistent across tutorials.

Crowd experts then create tasks to offer alongside their tutorial steps. Each step can have one or more tasks. Tasks include at least a description and price, but can also have examples of previous results, a custom email for responding to requests, and an estimated time to completion. Ink offers templates for common tasks (Table 2).

The expert can use Ink to create submission requirements for tasks. For example, in the tutorial "Write your Statement of Purpose for graduate school applications," the expert offered to read students' final essays and give them feedback. However, she required that they complete a checklist to verify their preparation before they submit. The combination of creating the tutorial and defining the tasks and requirements gives expert workers more agency to manage their work.

4.2 Task Assignment Keeps Tutorial Experts Available

Ink assigns all incoming tasks to the expert who created the tutorial. However, Ink workers are often experts who have other jobs and obligations, so they may not always be free to work. They

may also want to hire others to help manage the demand. Ink enables experts to add other workers to the queue for incoming tasks. For example, the writing expert who created the SoP tutorial could not work for a week, so she recruited one of her co-workers to take over temporarily. When she was back, she used a worker queue to continue assigning tasks to workers.

The queue functions by forwarding each new task to the expert at the top of the queue and asks them to confirm their availability. If they do not respond within 24 hours, Ink emails the next person. If they do accept the task, Ink assigns the task to that expert by sending a second email to both the user and the expert. Experts can also manage their workload via the queue, for example, by pre-accepting two tasks per day.

4.3 Hosting and Sharing to Bootstrap Web Traffic

Ink aims to help crowd experts kickstart their visitor base and to ease any technical burden with web hosting. When a worker submits the tutorial, it is hosted with a unique URL on the platform like any other CMS such as Medium or Wordpress.

Initially, we used our university's name to help the tutorials bootstrap credibility on the web and Ink began by hosting all tutorials at ink.stanford.edu. Even though the text and workers' bios said otherwise, users might have interpreted this to mean that our university was offering these services. In order to prevent this confusion, once the website had gathered reviews we changed Ink to a standalone website in October 2016, this was during the final two months of our evaluation. Specifically, we made the following changes: We changed the URL to inktutorials.org. We also added the following language under all task descriptions.

Ink is a research project at Stanford HCI. We have interviewed the experts who perform work on the platform and guarantee their proficiency. These experts are in no way affiliated with our university nor can they provide insider information about admissions to any university.

Therefore, we prevented any confusion about the expert crowd workers on the platform. All of the users who we interviewed in the next section had used Ink after this change in hosting.

4.4 Feedback and Reputation Help Verify Credibility

Ink hosts a reputation system consisting of star ratings and text reviews. Users who use the tasks are invited to rate their experience through a follow-up email. Ink embeds the star rating and the reviews on the front page for each tutorial. For experts, this reputation system is a way to signal the quality of their work. For visitors, it provides a more familiar reputation metaphor and helps ground their evaluation of the page's credibility. As a pragmatic bootstrapping step for the platform, we include a visible guarantee from our research group stating that we have interviewed all experts who host pages on Ink and verified that they can produce high quality results. If they fail to do so, our group will hire another expert to complete the task at no additional charge for the user. In practice, we only needed to do this twice for over 190 completed tasks. As Ink grows and gathers its own reputation, we expect that this guarantee could be hosted by Ink itself.

5 EVALUATION

Ink embodies the design hypothesis that expert crowd workers can embed their services within paths that users already travel. In this section, we investigated whether crowd experts can create effective tutorials that drive readership and usage, and whether users find the tutorials and paid help useful. We do so through a public deployment of Ink.

5.1 Field Deployment

We reached out to crowd workers on Upwork and worker forums for Amazon Mechanical Turk to create tutorials. We suggested tutorial subjects, but the experts had final power over the tutorials'

topics. We paid experts their hourly rate for creating this seed set of tutorials. Experts created tutorials that offered skills ranging from practicing technical interviews to feedback on academic writing and poster design:

Write a Statement of Purpose (SoP) for graduate school applications. 2,300 words. Author: An Upwork worker who is a PhD graduate in sociolinguistics and works at the Office of Scholarly Communication and Publishing at the University of Pittsburgh. Steps include developing a theme, outlining the main points, structuring an outline (with examples), and writing the main content. Tasks available: Feedback on the main points for $10, feedback on a full draft for $20.

Prepare for a technical interview. 600 words. Author: A former engineer and interviewer at a major software engineering company. The tutorial walks visitors through relevant readings, sources of practice problems, and the steps to follow after being asked a question. Task available: A practice technical interview for $20.

Write your TOEFL essay. 2,300 words. Author: An Upwork worker who holds a degree in teaching English as a second language. The tutorial explains the different types of essays (integrated and independent) and shares tips, then walks students through the steps of brainstorming, writing, and proofreading. Tasks available: Send your essay and get a score as well as feedback from a professional for $5.

Complete your first 5,000 tasks on Amazon Mechanical Turk. 550 words. Author: A long-time Turker. Steps include introducing tools that Turkers use and pointers to easy HITs to get started. Tasks available: Ask questions about qualifications, tools, or tasks from a professional Turker for 50 cents.

Prepare a 15-minute presentation. 550 words. Author: An artist and graphic designer who works on Upwork. Steps include preparing the main message; listing main topics and breaking them into sections; framing topics into intro, body, and conclusion; completing the content; and visual design. Tasks available: Visual design of slides by a professional graphic designer for $30, feedback on a recording from a professional speaking tutor for $20.

5.2 Method

We deployed the tutorials on the web. As the tutorials gathered users, we collected information about their experiences in two ways. First, we sought to learn about the general use cases for Ink, so we needed a method with a high response rate. We sent users a short survey with the following two questions: We asked them to rate their experience using Ink's five-star rating system and asked whether they had ever hired someone online for a task before. Second, to gain an in depth understanding of people's experience with Ink we contacted 30 people who had submitted task requests during November 2016 and requested to interview them. All of the users who we interviewed had used Ink after we had changed the URL to remove the Stanford domain. We conducted semi-structured interviews with 16 users who accepted our request. We interviewed participants on the phone whenever possible; however, due to the difficulties of arranging international phone calls and scheduling conflicts, we interviewed seven participants using text chat.

We consolidated users' responses and analyzed them according to the following themes: Reasons for using ink, prior experiences purchasing online services, difficulty of requesting tasks through Ink tutorials, trust in online services, and system design. We evaluate usage of Ink with a combination of qualitative data from our interviews and surveys and quantitative web log data.

Additionally, we sought to understand the effect of Ink's tutorial-embedded tasks in comparison to asking users to utilize a crowdsourcing marketplace. Was the major challenge for Ink users that

Table 2. Usage Results of the Tutorials from November 2015 to November 2016

Tutorial	Write a statement of purpose for graduate school applications	Create your academic curriculum vitae	Edit your paper before submission
Pageviews	8,000	509	714
Tasks completed	150	34	6
Avg. time on page	4:57	2:26	2:24
Main medium	Organic search (57.7%)	Direct (50.6%)	Direct (39.5%)

they did not know where they could go to get help? Or was the integrated task, with requirements and price, essential to its adoption?

To answer this question, we performed a one-month field experiment where users were randomly assigned to either the Ink interface or to a control interface where Ink tasks were replaced with a link to the expert's Upwork profile accompanied with the following text: "If you need more help after reading this tutorial you can hire an expert here: [link to Upwork]." We ran this experiment on the most popular tutorial: How to write your SoP for graduate school. The control condition provided a link to the tutorial author who had a public Upwork profile with high reputation (95% job success and over 2,000 hours experience). We measured the number of requests in each condition and performed a chi-square test of independence between the two conditions.

Separately, we sought to understand crowd experts' experience with Ink. Throughout the design and deployment of Ink, we worked closely with these experts as they created tutorials and gave feedback on the interface. After the tutorials were deployed and had gained usage, we engaged in email conversations with three experts who had posted tutorials on Ink and asked about their experience with Ink.

6 RESULTS

We focus our analysis on all 10 Ink tutorials and their activity over a year from November 2015 to November 2016. Table 2 shows usage on the three tutorials that received task requests during this period. Three other tutorials also received task requests but not in the period of our evaluation. During our evaluation, Ink tutorials received 25,000 pageviews from 7,000 unique visitors and 168 users paid an average of $20.30 (s.d. = $16.70) for 190 tasks.

One of the tutorials received a majority of the views and task requests, this is in line with known trends of online content creation that often follows an exponential curve (Guo et al. 2009). Closer attention to the main mediums for each tutorial suggests that the SoP tutorial showed up in more organic search results. This may be due to there being less online content on this topic than the other tutorials. Additionally, while many people viewed the tutorials, not all asked for direct help from experts. This is to be expected, as the tutorials are helpful standalone resources even without the crowd worker's direct services.

Of those who hired Ink workers, 115 responded to our follow-up survey for rating feedback, producing ratings that averaged 4.74 out of 5 (s.d. = 0.43).

All of the 16 participants (6 female, 10 male) who we interviewed had submitted a task request on Ink's most popular tutorial on writing a statement of purpose, and 15 had completed the transaction and received the service. These participants' ages ranged from 21 to 33 years old (M = 25) and they lived in Italy (1), South Korea (1), China (1), India (3), and the US (10). They rated the service that they had received 4.7/5 (s.d. = 0.5) in the post-survey.

In order to better understand Ink's common use mode, we begin with one case study.

6.1 Case Study: Statement of Purpose Tutorial

"Sue," a PhD graduate, who works at the University of Pittsburgh, spends some of her free time doing academic writing tasks on Upwork. We contacted her in the fall, around the time of PhD applications, and suggested that she create an Ink tutorial to help applicants write their SoP for graduate school applications. Relying on her own experiences with students and some internet research, she created a tutorial that guides students through the following steps: (1) Theme, (2) Main Points, (3) Structure, (4) Content. Sue steps in to offer help at two points. In the second step (Main Points), the student can upload their list of main points and get feedback on which points to highlight or remove from their final statement of purpose. At the final step (Content), the student can upload their essay and get feedback from Sue.

One student, "Eva," found the tutorial through a Google search. She decided to use the feedback task on the last step to get feedback, so she submitted her draft:

> Here is the SoP [statement of purpose] that I would like reviewed. I am applying to media and cultural studies doctoral programs.

Ink sent the following email to both of them:

> Hi, Sue [...] will be giving you feedback [...]. Please Paypal $20 to her at [...]. She will send you your feedback by Tuesday, December 9th.

Sue left comments and feedback on Eva's essay and sent it to her with some personal notes on her application. Eva responded:

> Thank you for your quick response time and your comments! The feedback was wonderful; I've implemented all of your suggestions. From a broader perspective, were there any areas that should be cut? Added? I was a little concerned that I may have made everything overly complicated or repetitive.

Sue responded to her questions with additional advice.

6.2 Usage Themes

In this section, we report on themes that emerged from qualitative data from our interviews with users, quantitative data from Ink's deployment, and results of our post-survey.

Ink reduces friction including both the information barrier and cost for hiring an expert. Ink addresses the information barrier for users to access expert crowd help and helps them identify opportunities. Additionally, all of our interview participants found getting help through Ink straightforward because instead of negotiating a hiring contract, Ink frames the process as submitting a request. This is a procedure that regular web users are familiar with:

> The website was very nice and has very good content. It gives clear directions on how I could submit my SoP. It was really helpful [...]. It was not difficult to determine steps that were to be followed. I simply went through the page and sent an email with my details.

We also found that many users valued the low financial cost of submitting a task because it reduced the risk of submitting a request and receiving low quality work. A total of 11 out of 16 participants in our interview mentioned the low cost and ease of use as the main reasons why they eventually decided to use Ink:

> It wasn't very expensive, I was just taking a chance [...] I didn't think that I had a need until I saw this and I thought I should try it

Therefore, Ink supports users to find low cost, context specific connections to experts who can help them. Because of their low cost, these types of connections are formed with greater frequency than on generic labor platforms. An initial low cost connection helps facilitate trust and increase the possibility of creating longer-term relationships.

A specific, context dependent need brought users to ink. One pattern repeatedly came up in our interviews when we asked participants what had brought them to Ink. All users said that they needed help with their statement of purpose for graduate school applications. This finding supports our initial hypothesis that people seek information in a narrow, context specific information space related to their need at that moment. On the other hand, generic labor platforms (e.g., Upwork and Fiverr) assume that a user enters seeking to hire help for a task at hand. This framing did not meet the needs of Ink users who often found the tutorial through a web search for information on a specific need. One participant who had found Ink's tutorial through a Google search said:

> I was looking to write a statement of purpose and I just typed in [a search engine...] I was really just looking for examples maybe of other people's and it was maybe the 3rd or 4th hit, so I clicked on it and I was like oh look, there's a service that lets you submit your draft and get it reviewed and it thought me what to write so that's what got me into it

Ink users would likely not have otherwise accessed expert crowd help. Of our 16 interview participants only 4 had hired someone online for a task before. This finding from our interviews was further supported through Ink's post-survey. Only 20% of 69 Ink users who responded to our post survey had previously sought paid expert help online. Following their regular practices, we expect that these users would not have hired a crowd worker for the task that they submitted to Ink either. Therefore, Ink connected users who would likely not have otherwise had access to crowd services to experts who could help them.

Users found the tutorials useful and engaged with the content. Users found the content and step-by-step structure of the tutorials useful as a guide to achieving their goal. Of 16 interview participants, 15 told us that they read the tutorial, and either followed it directly or used it to retrospectively analyze their essay:

> It was structured well, the website, and walking through the steps of how to write an effective essay and I thought it was very informative and clearly written.

The tutorial also served as a means to add context to the service that was offered. When we asked a participant who had heard of online labor markets why he did not approach one of those websites to get help on his essay he said:

> I think I could have probably found someone on there [Upwork] but this was streamlined specifically for the specific task. Like, [on] Upwork, I know they have people who could review it but this was like statement of purpose specifically and I read the materials about how to like write the statement of purpose so I was like ok these guys are experts in this specific thing it's not just like hiring a general TaskRabbit.

Web analytics also suggest that the tutorials were engaging. On average each web session lasted over 1 minute, enough time for users who arrived at Ink tutorials to read the content. Tutorials

varied in their popularity. For example, the tutorial "How to write your SoP for graduate school applications" was viewed 8,000 times, with an average time on page of 5 minutes. The tutorial on "Edit you paper before submission" was viewed 714 times with an average time on page of over 2 minutes.

In total, 30% of users continued their conversations with experts after the first task. Once they had established a working relationship with an expert, 30% of users contacted them directly for questions or future needs. For example, some students who had gotten help with writing their SoP hired the same expert again later on to get help with other writing tasks.

Users benefited from the familiarity and trust that they had built with the online expert. For instance, after receiving feedback on her statement of purpose, one student requested feedback from the same expert on three other essays that she was working on, she explained her reason:

> I liked her feedback, and the continuity of having one person review my essays.

These connections sometimes took the form of hands-on mentorship. For instance, one student who had done a practice technical interview asked our expert to help her friend become familiar with computer science interviews. The expert contacted us:

> Alice, one of the people I interviewed, wants to do another Ink session. Her friend is also interested but has never interviewed before. Can she attend my session with Alice before scheduling her own?

Users' trust in ink was enhanced by the tutorial quality, reviews, and research institution name. In the absence of specialized payment infrastructure, users of our system were required to Paypal the full fee (e.g., $20 for the SoP tutorial) directly to the expert crowd worker in advance of receiving the service. We asked interview participants about the reasons that they decided to trust the platform. Of 16 participants, 11 mentioned that they valued the reviews on the website as an indicator of quality of service and 8 mentioned that having our research group's name on the website was their major reason for trusting that the website was not scam and they could pay up-front for tasks. In deploying Ink, we initially relied on our university affiliation to bootstrap trust, but changed to an independent URL after the website had gathered some usage. We found associating with a credible institution, a good method for bootstrapping such systems.

Additionally, eight participants mentioned the quality of the tutorial as an indicator of high quality work and the reason that they trusted the website:

> The tips and examples that were given on the website seemed very in-line with what I was looking for [...]. Seeing that and seeing that there was already information being provided made me think that then I would receive additional [high] quality information for the paid service.

Organic web search was the main driver of traffic to the tutorials; people also shared the tutorial with friends. Most interview participants (9/16) told us that they had found Ink via a web search for content to help them in writing their statement of purpose. The self-contained nature of the tutorials also made it easy for people to share with their friends. Four of our interview participants had been referred to Ink by a friend who had used the system before. Three had found Ink from other online resources such as forums and subReddits.

This result was further confirmed by the website's analytics, from November 2015 to 2016, 51% of all Ink sessions originated from organic web search, 26% direct, 13% referral, and 10% social. The most-viewed tutorial, "How to write your SoP for graduate school applications," had a similar distribution across channels, i.e., 58% organic search and 30% direct.

Half of ink tutorials received no requests. Of nine tutorials that experts created on Ink, four received no requests and one received only one request. Therefore, for Ink experts gathering attention to their tutorials was a challenge. Early on, crowd experts used Google Adwords, mailing lists, and social media to attract users to Ink tutorials. This approach was not very successful. While it increased viewers, it did not increase usage of crowd workers' services. Organic attention (e.g., through web search) proved necessary for users to have a higher level of commitment. Furthermore, Ink tutorials that covered subjects and offered help on areas already prevalent on the web were less successful at attracting attention. For example, the tutorial offering to help with web design struggled to gain traction, while the tutorial for writing a statement of purpose was very successful.

Therefore, experts were most successful when they found a sweet spot between domains that were popular enough to gain traction while not too prevalent on the web. Experts may also find it beneficial to start with domains that are long-tail information needs and move into new domains as they gather reputation.

6.3 Field Experiment Results

We analyzed the number of requests from the field experiment where we manipulated whether Ink tasks were embedded in the tutorial or redirected to Upwork. While the Ink condition received 32 requests during the month of the field experiment, only one request was submitted in the control condition. A chi-square test of independence shows that the relationship between these variables is significant $\chi^2(1, N = 33) = 29.1, p < 0.05$. Thus, users were not only struggling to find experts to hire, but found the process of approaching the marketplace to request tasks onerous enough to avoid.

6.4 Reflections from Expert Crowd Workers

We interviewed three experts who had created Ink tutorials. These experts were all Upwork workers with high ratings and a cumulative work experience of 2,770 hours on Upwork. We consolidated experts' responses and analyzed them according to the following predetermined themes framed by our research question: prior experience with expert crowd work, authoring tutorials, creating tasks, power relations in online labor markets, relations with users, and system design feedback. The following themes emerged on these expert crowd workers' experience with Ink.

Workers reported that defining their own tasks granted them greater feelings of control. Crowd workers felt that they had more control over their tasks with Ink's model:

> I really like the idea of creating the tasks myself. It gives me a sense of control over the work that I would be potentially doing in the future. [...On Upwork] you might not know exactly what your work will look like, depending on your client's wishes along the way. What sounds like a simple task might become endlessly complicated, or the client might want something from you that they didn't put into the job ad. With your style of work, it's almost like the roles are reversed.

Therefore, by creating their own tasks, expert crowd workers may be able to set expectations for users and take more of a managerial role. This finding is important because these workers have extensive experience in a crowdsourcing marketplace; this "sense of control" contrasts with the unbalanced power relations that are common to crowdsourcing systems (Irani and Silberman 2013; Salehi et al. 2015). Given the limited number of workers in our deployment, however, more research is needed to generalize this finding to a larger ecosystem.

Creating tutorials gives experts more flexibility to be creative. A major challenge in current expert crowdsourcing systems is that users post the jobs, and workers must choose to take them or leave them. On Upwork, for example, users have specific requests and workers must pay close attention to work within those requirements. While creating Ink's tutorials, however, experts experimented with new, creative domains. One expert said:

> I really enjoy the creative freedom of creating the tutorials; [...] I can rely on resources that I know of and my own experience to create a new thing that might actually help someone out there.

In the next section, we discuss the potential of re-designing hiring surfaces for crowd work that afford more agency to workers and some open questions that this creates for future work.

7 DISCUSSION AND FUTURE WORK

In this article, we have shown that users face barriers in hiring expert crowd workers and that hiring surfaces can be designed to mitigate those challenges by giving crowd workers more agency to frame and manage their work. Our goal was to reduce the friction that regular end users experience in hiring expert crowd services. In doing so we envision a future where by receiving help from online experts through written guides and personalized tasks, users will be able to reach goals that would have otherwise been inaccessible. In this section, we discuss our system and its limitations. We also outline challenges and open questions for future work in this area.

Ink's model is best suited for goals that can be formed into step-by-step items within a tutorial. Ink relies on users recognizing an opportunity within context and utilizing it; therefore, Ink tasks need to be well defined and repeatable. Ink tasks also need to be fairly common in order to make the effort of creating the tutorial and publishing it worthwhile. For example, future Ink tutorials can be utilized to help users file their taxes, write a paper, or decide on a purchase. Goals that are unclear at first and whose needs are uncommon may be better suited for the traditional market model where workers with relevant skills bid on tasks and engage in back-and-forth communication with the user to reach mutual understanding (Salehi et al. 2017).

One challenge of this style of creative work is building confidence. We found that experts were hesitant to spend time on a tutorial if they were not sure that anyone was going to use it. Seeding a tutorial with basic content may be a low-risk way for workers to test the waters: if the tutorial gathers initial pageviews, it may be worth investing more time into. Future systems can design mechanisms to lower the risk by encouraging experts to take the first step, and inviting them to work more if they find traction.

7.1 Sources of Friction for Users

In this article, we have proposed that threading pre-defined tasks within users' regular information seeking habits reduces friction for their adoption of crowd work services. In this section, we discuss the sources of this friction.

The first prominent barrier is recall. Human behavior is often automatic and unreflective (Weiss and Ilgen 1985; Kahneman 2011), so even if someone knows about paid crowd work, they may not recall it as an option when faced with a new task. Searching online for web content, in contrast, is a deeply-ingrained routine behavior for information tasks. Ink relies on this insight to reduce recall as a source of friction.

Information about a worker's skills and trust in the quality of the service is another potential source of friction for users. Current platforms use review systems similar to product reviews to help resolve this inefficiency. However, these systems are limited because novice users struggle to review experts and a combination of avoiding social awkwardness and a lack of incentives result

in reviews that are inflated and uninformative (Whiting et al. 2017). Daemo's crowd guilds offer a potential solution in which workers review their peers (Whiting et al. 2017). Another example, Boomerang, makes the review system incentive compatible by ensuring that high reviewed workers get early access to a requester's tasks (Gaikwad et al. 2016).

Up & Go[2] a second-level co-op platform for home services takes a combination of approaches to maintain trust in the quality of their services. Up & Go uses a rating and review system that is applied to the cooperative business rather than an individual worker. A quality satisfaction policy hosted by the platform ensures that if a customer is not satisfied with the service they received the cooperative will return within 72 hours to finish the parts of the job that were not done to satisfaction. Furthermore, to ensure high quality work and to maintain trust, each member first goes through an application process and candidacy period in which they are evaluated. Finally, as a co-owner of the co-op each worker is "invested in the success of the business and looking to gain loyal customers and remain accountable to their fellow cooperative members, customers, and community." Future research will further develop models that increase trust between users and workers.

7.2 Crowd Worker Agency and the Challenge of Scale

With Ink, we have proposed an alternative design of expert crowd work systems with the goal of reducing friction for end users. Toward this goal, we rely on the labor and creativity of expert crowd workers who make Ink function, while granting these workers greater flexibility and control over their work than traditional crowd work models. We argue that such design decisions may have a positive impact on the experience of crowd work in the future.

In this article, we have shown Ink's viability in a small-scale longitudinal deployment. We have shown that expert crowd workers can create tutorials and offer their own services, benefiting over two hundred real-world users, most of whom would not have otherwise approached an online labor market. The complexities that will emerge as Ink scales can be analyzed at two levels, i.e., having more work than the tutorial author can manage, and growth in the number of tutorials.

The first challenge occurred during our deployment where the SoP tutorial received many requests close to application deadlines. In our case, the tutorial author recruited one of her coworkers to take on some of the workload. Likewise, prior work has proposed subcontracting crowd work and defining managerial roles for workers (Morris et al. 2017). In our case, the tutorial author used a basic queue to divide tasks between the two workers. Later, she expanded the queue to five workers including herself. Future research can design other models. For instance, the tutorial author may want to review finished work before it is shipped to the user.

The second challenge occurs when the number of tutorials scale and there may be more than one tutorial on a given subject. We leave this problem largely to future work, but suggest two scenarios: First, the platform may allow for authors to submit as many tutorials as they want to, some covering the same topic. This pattern of open user-generated content is similar to Etsy or Youtube, where producers create a range of products or videos that compete. In this case, search engines can be developed to bring attention to the various resources available for a user query. Alternatively, workers on the platform may decide to limit repetitions and take on an approach similar to Wikipedia, where each topic is covered by only one article. This approach would require platform level governance by users to make decisions on adding new tutorials, and managing current ones.

In both cases scalable models need to be developed to assure that worker agency does not get restricted with growth and that design happens in collaboration with workers and toward their

[2]https://www.upandgo.coop.

best interests (Irani and Silberman 2016). Some examples of work in this direction are Daemo, an opensource crowd work platform governed by workers and requesters that is governed by a shared constitution[3] and Up & Go[4] a second level co-op platform of worker-owned co-ops for home services. Both of these examples are in early stages and further research will study scalable mechanisms for their growth.

8 CONCLUSION

In this article, we have focused on resolving friction for utilizing crowd expertise by users. We have addressed this problem by creating a system, Ink, that experts can use to thread tasks into users' paths via web tutorials. In a public deployment, over 200 web users hired crowd experts from these tutorials. We propose that hiring surfaces for crowd work can be designed to give workers more agency over their work and discuss open challenges toward that goal.

ACKNOWLEDGMENTS

The authors would like to thank the expert crowd workers who contributed to this project.

REFERENCES

Lada A. Adamic, Jun Zhang, Eytan Bakshy, and Mark S. Ackerman. 2008. Knowledge sharing and Yahoo answers: Everyone knows something. In *Proceedings of the 17th International Conference on World Wide Web*. ACM, 665–674.

Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A word processor with a crowd inside. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*. ACM, 313–322.

Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samual White, and Tom Yeh. 2010. VizWiz: Nearly real-time answers to visual questions. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*. ACM, 333–342.

Joel Brandt, Philip J. Guo, Joel Lewenstein, Mira Dontcheva, and Scott R. Klemmer. 2009. Two studies of opportunistic programming: Interleaving web foraging, learning, and writing code. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1589–1598.

Andrea Bunt, Patrick Dubois, Ben Lafreniere, Michael A. Terry, and David T. Cormack. 2014. TaggedComments: Promoting and integrating user comments in online application tutorials. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 4037–4046.

Yan Chen, Teck-Hua Ho, and Yong-Mi Kim. 2010. Knowledge market design: A field experiment at google answers. *Journal of Public Economic Theory* 12, 4 (2010), 641–664.

Snehalkumar Neil S. Gaikwad, Durim Morina, Adam Ginzberg, Catherine Mullings, Shirish Goyal, Dilrukshi Gamage, Christopher Diemert, Mathias Burton, Sharon Zhou, Mark Whiting, and others. 2016. Boomerang: Rebounding the consequences of reputation feedback on crowdsourcing platforms. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 625–637.

Mary L. Gray, Siddharth Suri, Syed Shoaib Ali, and Deepti Kulkarni. 2016. The crowd is a collaborative network. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 134–147.

Lei Guo, Enhua Tan, Songqing Chen, Xiaodong Zhang, and Yihong Eric Zhao. 2009. Analyzing patterns of user content generation in online social networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 369–378.

Lilly C. Irani and M. Silberman. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 611–620.

Lilly C. Irani and M. Silberman. 2016. Stories we tell about labor: Turkopticon and the trouble with design. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4573–4586.

Sheena S. Iyengar and Mark R. Lepper. 2000. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology* 79, 6 (2000), 995.

Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Macmillan.

[3] https://crowdresearch.github.io/open-gov/constitution/2016/11/21/constitution.html.

[4] https://colab.coop/work/upandgo/.

Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*. ACM, 1301–1318.

Ben Lafreniere, Andrea Bunt, Matthew Lount, and Michael A. Terry. 2013. Understanding the roles and uses of web tutorials. In *ICWSM*.

Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1603–1612.

Lena Mamykina, Bella Manoim, Manas Mittal, George Hripcsak, and Björn Hartmann. 2011. Design lessons from the fastest Q&A site in the west. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2857–2866.

David Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a turker. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 224–235.

Meredith Ringel Morris, Jeffrey P. Bigham, Robin Brewer, Jonathan Bragg, Anand Kulkarni, Jessie Li, and Saiph Savage. 2017. Subcontracting microwork. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM.

Daniela Retelny, Sébastien Robaszkiewicz, Alexandra To, Walter S. Lasecki, Jay Patel, Negar Rahmati, Tulsee Doshi, Melissa Valentine, and Michael S. Bernstein. 2014. Expert crowdsourcing with flash teams. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*. ACM, 75–85.

Niloufar Salehi, Lilly C. Irani, Michael S. Bernstein, Ali Alkhatib, Eva Ogbe, Kristy Milland, and others. 2015. We are dynamo: Overcoming stalling and friction in collective action for crowd workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1621–1630.

Niloufar Salehi, Jaime Teevan, Shamsi Iqbal, and Ece Kamar. 2017. Communicating context to the crowd for complex writing tasks. In *Proceedings of the ACM 2017 Conference on Computer Supported Cooperative Work*. ACM.

Ryo Suzuki, Niloufar Salehi, Michelle S. Lam, Juan C. Marroquin, and Michael S. Bernstein. 2016. Atelier: Repurposing expert crowdsourcing tasks as micro-internships. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2645–2656.

Amos Tversky and Daniel Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5, 4 (1992), 297–323.

Benjamin O. Weeks. 2004. The paradox of choice: Why more is less. *The Academy of Management Executive* 18, 4 (2004), 170–171.

Howard M. Weiss and Daniel R. Ilgen. 1985. Routinized behavior in organizations. *Journal of Behavioral Economics* 14, 1 (1985), 57–67.

Mark E. Whiting, Dilrukshi Gamage, Snehalkumar Neil S. Gaikwad, Aaron Gilbee, Shirish Goyal, Alipta Ballav, Dinesh Majeti, Nalin Chhibber, Angela Richmond-Fuller, Freddie Vargus, Tejas Seshadri Sarma, Varshine Chandrakanthan, Teogenes Moura, Mohamed Hashim Salih, Gabriel Bayomi Tinoco Kalejaiye, Adam Ginzberg, Catherine A. Mullings, Yoni Dayan, Kristy Milland, Henrique Orefice, Jeff Regino, Sayna Parsi, Kunz Mainali, Vibhor Sehgal, Sekandar Matin, Akshansh Sinha, Rajan Vaish, and Michael S. Bernstein. 2017. Crowd guilds: Worker-led reputation and feedback on crowdsourcing platforms. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 1902–1913.

Received February 2017; revised November 2017; accepted December 2017