

Crowd science user contribution patterns and their implications

Henry Sauermann^{a,1} and Chiara Franzoni^b

^aGeorgia Institute of Technology, Scheller College of Business, Atlanta, GA 30308; and ^bPolitecnico di Milano, School of Management, Milan 20133, Italy

Edited by Adrian E. Raftery, University of Washington, Seattle, WA, and approved December 11, 2014 (received for review May 13, 2014)

Scientific research performed with the involvement of the broader public (the crowd) attracts increasing attention from scientists and policy makers. A key premise is that project organizers may be able to draw on underused human resources to advance research at relatively low cost. Despite a growing number of examples, systematic research on the effort contributions volunteers are willing to make to crowd science projects is lacking. Analyzing data on seven different projects, we quantify the financial value volunteers can bring by comparing their unpaid contributions with counterfactual costs in traditional or online labor markets. The volume of total contributions is substantial, although some projects are much more successful in attracting effort than others. Moreover, contributions received by projects are very uneven across time—a tendency toward declining activity is interrupted by spikes typically resulting from outreach efforts or media attention. Analyzing user-level data, we find that most contributors participate only once and with little effort, leaving a relatively small share of users who return responsible for most of the work. Although top contributor status is earned primarily through higher levels of effort, top contributors also tend to work faster. This speed advantage develops over multiple sessions, suggesting that it reflects learning rather than inherent differences in skills. Our findings inform recent discussions about potential benefits from crowd science, suggest that involving the crowd may be more effective for some kinds of projects than others, provide guidance for project managers, and raise important questions for future research.

crowd science | citizen science | crowdsourcing | dynamics | effort valuation

A growing number of scientific projects involve members of the general population (the “crowd”) in research. In most cases, unpaid participants (also called citizen scientists) perform relatively simple tasks such as data collection, image coding, or the transcription of documents (1–3). Sometimes they directly make important discoveries, such as new classes of galaxies in the project Galaxy Zoo or the structure of proteins relevant for the transmission of HIV in the project Foldit (4–6). Such discoveries may be explicit goals of a project, but they have also occurred serendipitously as members of the crowd were working on their primary tasks. Although computer technologies can be used for an increasing number of research activities, human intelligence remains important for many tasks such as the coding of fuzzy images, optimization with high degrees of freedom, or the identification of unexpected patterns (7–9).

Recent discussions suggest at least six types of benefits from involving the crowd in the production of scientific research (2, 3, 10, 11). First, members of the crowd typically contribute because of intrinsic or social motivations rather than for financial compensation (2, 12), potentially allowing project organizers to lower the cost of labor inputs compared with traditional employment relationships. Second, crowd science projects provide important speed advantages to the extent that a large number of contributors work in parallel, shortening the time required to perform a fixed amount of work. Third, by “broadcasting” their needs to a large number of potential contributors, projects can gain access

to relatively rare skills and knowledge, including those that are not typically part of scientific training (10). Fourth, projects that require creative ideas and novel approaches typically benefit from rich and diverse knowledge inputs (13), and involving a larger crowd of individuals with diverse competences and experiences is more likely to provide access to such inputs. Fifth, crowd science projects can involve contributors across time and geographic space, allowing them to increase coverage that is particularly important for observational studies (14). Finally, in addition to potential impacts on productivity, involving the general public in research may also yield benefits for science education and advocacy (1, 15, 16).

In light of these potential benefits, crowd science is receiving increasing attention within and outside the scientific community. For example, the National Institutes of Health (NIH) is discussing the creation of a common fund program for citizen science (17), and the US Federal Government highlights the crowdsourcing of science as a key element in its Open Government National Action Plan (18). Google recently gave a \$1.8 million Global Impact Award to the Zooniverse crowd science platform (19), and The White House honored four citizen science leaders with a “Champions of Change” award (20). Similarly, case studies of crowd science projects as well as policy articles on the subject have been published in major scientific journals such as PNAS, *Science*, and *Nature* (6, 7, 21).

Despite the growing attention, our empirical understanding of crowd science is limited. The existing literature has demonstrated the feasibility of crowd-based approaches using case examples and has explored organizational innovations to increase efficiency or expand the range of tasks in which volunteers can participate (6, 7, 22, 23). Even though project success depends critically on the labor inputs of the crowd, there

Significance

Involving the public in research may provide considerable benefits for the progress of science. However, the sustainability of “crowd science” approaches depends on the degree to which members of the public are interested and provide continued labor inputs. We describe and compare contribution patterns in multiple projects using a range of measures. We show that effort contributions can be significant in magnitude and speed, but we also identify several challenges. In addition, we explore some of the underlying dynamics and mechanisms. As such, we provide quantitative evidence that is useful for scientists who consider adopting crowd science approaches and for scholars studying crowd-based knowledge production. Our results also inform current policy discussions regarding the organization of scientific research.

Author contributions: H.S. and C.F. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

¹To whom correspondence should be addressed. Email: henry.sauermann@scheller.gatech.edu.

is little quantitative evidence on the labor flows to projects and on the degree to which users are willing to contribute time and effort to projects on a sustainable basis. Patterns of volunteer contributions have been explored in other contexts such as Open Source Software (OSS) and Wikipedia (24–27), yet crowd science projects differ significantly with respect to their organization, the types of rewards offered to participants, and the nature of the tasks that are crowdsourced (2, 7, 12, 28). Moreover, little is known about how contribution patterns differ (or not) across different projects and about the micromechanisms that shape the observed macro patterns.

We examine contribution patterns in seven different crowd science projects hosted on the platform [Zooniverse.org](#), currently the largest aggregator of crowd science projects in various disciplines. Drawing on over 12 million daily observations of users, we quantify and compare effort contributions at the level of the projects but also provide insights into individuals’ micro-level participation patterns. We then discuss implications for projects’ ability to garner the above-mentioned benefits of involving the crowd. Our findings are of interest not only to scholars studying crowd science or human computation; they also inform the broader scientific community and policy makers by providing insights into opportunities and challenges of a new organizational mode of conducting scientific research. Moreover, they point to a number of important questions for future research. Notwithstanding potential benefits for science education, our discussion will focus on the potential of crowd science to advance the production of scientific knowledge.

Results

We analyze data from the seven projects that were started on [Zooniverse.org](#) in 2010: Solar Stormwatch (SS), Galaxy Zoo Supernovae (GZS), Galaxy Zoo Hubble (GZH), Moon Zoo (MZ), Old Weather (OW), The Milkyway Project (MP), and Planet Hunters (PH). Although hosted on the same platform and using some shared infrastructure, these projects are run by different lead scientists, span different fields of science, and ask volunteers to perform different types of tasks (see [SI Text](#) and [Table S1](#) for background information). We use daily records of individual participants observed over the first 180 d of each project’s life. We use two measures of participants’ contributions: the time spent working and the number of classifications made on a given day. The generic term “classification” refers to a single unit of output in a particular project, e.g., the coding of one image, the tagging of a video, or the transcription of a handwritten record from the logbook of an old ship.

Large Supply of Labor at Low Cost. A total of 100,386 users participated at least once in a project in the first 180 d. However, there is considerable heterogeneity across projects (Fig. 1). The largest project (Planet Hunters) attracted 28,828 users in its first 180 d, whereas the smallest (Galaxy Zoo Supernovae) attracted 3,186. These users contributed a total of 129,540 h of unpaid labor, ranging from more than 54,000 h for the largest project to about 1,890 h for the smallest ([Table S2](#)). By showing the number of contributors and of hours contributed for each project, Fig. 1 also highlights differences across projects in the average amount of effort contributed per user, which ranges from roughly half an hour for Solar Stormwatch to more than 3.5 h for Old Weather. We discuss potential drivers of heterogeneity across projects in [SI Text](#).

We estimate the monetary value of these unpaid contributions using two complementary approaches ([SI Text](#)). The first approach is input based and values the hours contributed at the typical hourly wage of US undergraduate research assistants (\$12.00/h). Using this approach, the total contributions amount to \$1,554,474 and range from \$22,717 for the smallest project to \$654,130 for the largest, with an average of \$222,068 (Fig. 2). An alternative approach focuses on the outputs of the work and estimates the cost of the classifications on Amazon Mechanical




Fig. 1. For each of the seven projects, this figure shows the total number of users during the first 180 d of the project’s life (in thousands, *Left*, left axis), as well as the total number of hours contributed by those users (in thousands, *Right*, right axis).

Turk (AMT), a crowdsourcing platform on which paid workers perform microtasks that are very similar to those performed in our crowd science projects. Using this approach, projects received classifications valued at \$1,543,827, ranging from \$29,834 to \$485,254 per project, with an average of \$220,547. Note that these figures refer only to the first 180 d of each project’s life; contributions received over the whole duration of the projects (typically multiple years) are likely considerably larger. Although the two approaches yield very similar estimates of total value across all seven projects, the two estimates differ somewhat more for individual projects (Fig. 2). These differences in the valuation of contributions may reflect differences in hourly vs. output-based pay, i.e., it appears that AMT piece rates are not always proportional to the time required to perform a particular task. Overall, our estimates suggest that, although running crowd science projects is by no means costless (Zooniverse has an operating budget funded primarily through grants), drawing on unpaid volunteers who are driven by intrinsic or social motivations may indeed enable crowd science organizers to perform research that would strain or exceed most budgets if performed on a paid basis.

A Small Share of Contributors Makes a Large Share of the Contributions.

Analyses in other crowd-based settings such as Wikipedia and OSS development suggest that distributions of individuals’ contributions are highly skewed (24–26), and case evidence suggests uneven contributions also in crowd science (7, 29). However, comparative estimates for crowd science are lacking. Using an approach frequently used to study inequality of income (30), we plot the Lorenz curve for the distribution of individuals’ total classifications for each project in Fig. 3 and also compute the corresponding Gini coefficients ([SI Text](#)). The large area between the Lorenz curves and the 45° line (i.e., perfect equality) indicates that most of the contributions are supplied by a small share of volunteers in all seven projects. For example, the top 10% of contributors (top contributors) supplied between 71% and 88% of all contributions across projects, with an average of 79%. The Gini coefficients range from 0.77 to 0.91 (average, 0.85). These values are somewhat higher than those found in OSS development (Gini between 0.75 and 0.79) (26, 31, 32), but indicate less concentration than contributions to Wikipedia (0.92 or more) (24). Even among our seven projects, however, the distributions are significantly different from each other ([SI Text](#)). Thus, although crowd-based projects generally appear to be characterized by strong inequality in individuals’ contributions,




Fig. 2. (Left) Counterfactual cost of contributors' effort evaluated at an hourly rate of \$12. (Right) Counterfactual cost of the classifications made by contributors, evaluated using piece rates for similar microtasks on Amazon Mechanical Turk.

future work is needed to understand why the distributions of contributions are more skewed in some projects than in others (33).

Skewed distributions raise the question of who the top contributors are and what motivates their high levels of activity. Although our data do not allow us to explore the motivations of top contributors (for some fascinating qualitative accounts, see refs. 7 and 34), we can seek to further characterize them in terms of their activity in the project. A first possibility is that top contributors work faster than others. Using the average time taken per classification as a measure of speed, we find that the top 10% of contributors are indeed faster in six of the seven projects, although that advantage is relatively small (on average 14% less time per classification than the non-top contributors; Table S2). Interestingly, supplementary analyses suggest that top contributors do not typically have a speed advantage on their first day in a project, but their speed increases over subsequent visits, suggesting that higher speed results from learning rather than higher levels of innate skills or abilities (SI Text, Fig. S1, and Table S3). Because the speed advantage over non-top contributors is relatively small, however, top contributors' higher quantity of classifications primarily reflects more hours worked. This leads us to examine more closely the dynamics of individuals' effort.

Contributions Are Primarily Driven by Those Who Return for Multiple Days. Prior evidence from contexts such as OSS development or Wikipedia suggests that volunteer participation tends to be infrequent and of low intensity (25, 33). We find similar patterns in crowd science projects. First, most users do not return to a project for a second time, with the share of those who return ranging from 17% to 40% (average, 27%; Table S2 and Fig. S2). Second, when averaging the daily time spent across all active days for a given contributor, the mean ranges from 7.18 to 26.23 min across projects, indicating that visits tend to be quite short (Table S2 and Fig. S3).

To examine how total effort contributions depend on the intensity of effort at a given point in time vs. the frequency of participation over time, we distinguish four groups of contributors in each project: those that contribute only for 1 d and for a relatively short time on a given day (defined as below the 90th percentile in average time per active day), those that contribute only once but for a long time, those that contribute on multiple days but for a short time, and those that contribute on multiple days for a long time. Fig. 4A shows the share of users falling into each of the four groups and Fig. 4B shows the share of total

contributions made by each group of users. Although there is considerable heterogeneity across projects (both with respect to the distribution of the four groups and with respect to the share of effort contributed by each group), some general patterns emerge. In particular, the large majority of users contributed only on 1 d and for a relatively short time. Contributors working on multiple days and for a long time per day are a small minority of 4–7%. However, except for Solar Stormwatch, the latter group contributes the largest share of classifications. By distinguishing between the frequency of activity in a project and the average duration of effort on a given day, Fig. 4 also shows that the largest contributions are made primarily by users who return for multiple days (even if the effort per day is low) rather than by those who participate for a long time on a single day. Without users who return, projects would still have 73% of their members, but only receive 15% of the classifications.

Given that returning users are so important, we investigated their return behaviors in more detail by examining the distribution of active days in a project (Fig. S2), as well as the timing of active days. With respect to the latter, we find that even those users who return multiple times do so only after several days and with increasingly long breaks between visits. For example, among users with at least 7 active days (4.13% of the sample), the average break between active day 1 and 2 is 5.23 d and increases to 8.30 d between visits 6 and 7 (SI Text). Thus, participation frequency declines even for highly active users. Finally, we investigated which users are most likely to return to a project. Because we lack background information on individual users, we related return behavior to the time at which users initially joined the project (e.g., early vs. late in a project's life). Although several significant relationships emerge within particular projects, the strength and direction of these relationships differ across projects, providing no consistent overall picture (see SI Text and Table S4 for details).

Contributions Received by Projects Are Highly Volatile and Critically Depend on New Users. Having explored contribution dynamics at the level of individual contributors, we now examine how activity develops over time at the level of projects. Fig. 5 plots the average number of daily hours received by all projects over the first 180 d of their life (separate graphs for each project are shown in Fig. S4). Focusing first on the top line indicating the total volume




Fig. 3. Lorenz curves for the distribution of users' total classification counts in each project. The Lorenz curve indicates the cumulative share of classifications (y axis) made by a particular cumulative share of users (x axis). The stronger the curvature of the Lorenz curves, the stronger the inequality in contributions. For comparison, we also show the 45° line, which corresponds to total equality, i.e., all users contribute the same amount.

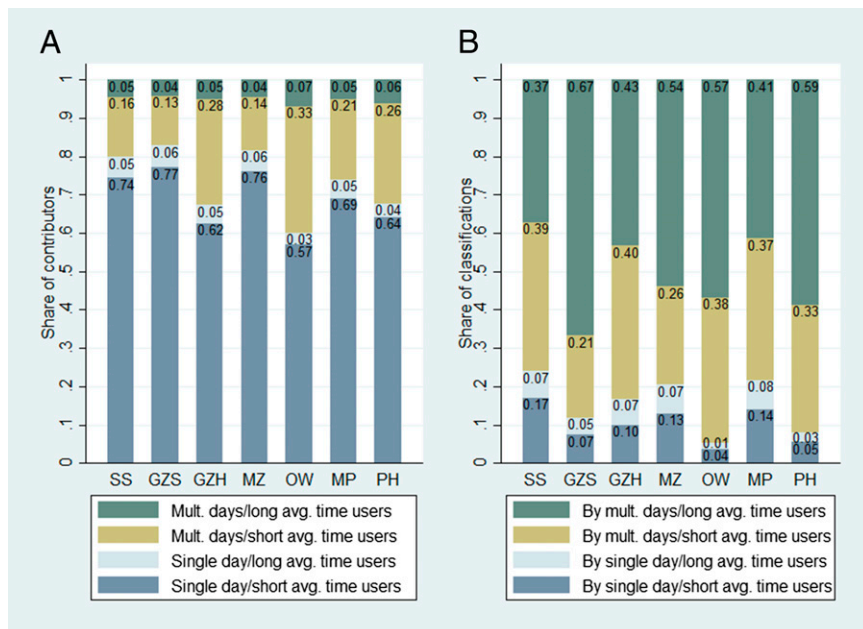


Fig. 4. (A) Share of four groups of users in each project. The four groups are defined by two dimensions: single-day participation vs. multiple-day participation and short average time spent per active day (below 90th percentile in project) vs. long average time (above 90th percentile). (B) Share of classifications made by each group of users.

of hours received, we find that Zooniverse projects attract the highest level of contributions early on in their life, followed by declining activity over time. Despite this general trend, daily contributions are very volatile with noticeable spikes. The front-loaded pattern of activity may reflect that the platform announces the launch of a new project to a large existing installed base of Zooniverse members. One might expect a different pattern in stand-alone projects, where knowledge about the project may diffuse more gradually among potential contributors. Although we lack detailed comparison data on stand-alone projects, a front-loaded pattern of user activity has also been reported from the project Phylo, and its organizers attribute this pattern to the (initial) novelty of the project and broad media coverage when the project started (9).

To gain a deeper understanding of the different sources of effort contributions over time, Fig. 5 divides the total number of hours received into the contributions made by three groups of users: original contributors, defined as participants who joined during the first 7 d of a project's life, new contributors, defined as those who joined within a rolling window of the last 7 d before the observation date, and a third residual group (i.e., those who joined after day 7 of a project's life, but more than 7 d before the observation date). Note that across projects, an average of 22% of users are original contributors; the remaining individuals start as new contributors when they join the platform but move into the middle category after 7 d. Fig. 5 shows that, although original users are (by definition) the sole contributors in the first days of the project, their contributions decline quickly, consistent with our earlier findings regarding the short duration of most users' activity in a project. Significant levels of effort are received from new users who continuously join throughout the observation period, but their effort does not compensate for the loss of effort from original users. As such, the daily amount of effort received by projects tends to decline over time. An analysis of cumulative effort contributions shows that over the 180-d observation window, original users are responsible for an average of 33% of total hours across the seven projects (Table S2 and Fig. S5).

Finally, Fig. 5 shows that some of the spikes in contributions reflect higher levels of activity from all types of users, i.e., projects attract effort from new users but also reactivate existing users (see also Fig. S6). Other spikes reflect primarily contributions from new users. Discussions with project organizers suggest that spikes can have a number of causes including

organizers' outreach efforts via email or social media, as well as attention from mainstream media. For example, the spike around day 120 coincides with an article featuring the project Planet Hunters in Time Magazine on April 14, 2011. We explore spikes in activity further in *SI Text*.

Discussion

We complement prior qualitative work on crowd science by characterizing and comparing contribution patterns in multiple crowd science projects using quantitative measures. Our analysis shows that projects can attract considerable volunteer effort that would be difficult and costly to procure through traditional or online labor markets. Using two different approaches, we estimate the average value of contributions received per project at more than \$200,000 over the first 180 d. However, some projects




Fig. 5. Area chart of average daily hours of effort received by the seven projects during the first 180 d of each project's life. Total effort divided into hours spent by original users (users who joined in the first 7 d of the project; *Bottom*), new users (users who joined in a rolling window of the last 7 d; *Top*), and the residual group (users who joined after day 7 of the project but >6 d before the observation day; *Middle*). Separate graphs for each project are shown in Fig. S4.

are much more successful in attracting contributions than others. Moreover, the volume of contributions shows high variability over time, with noticeable spikes that are typically driven by outreach efforts or media attention. Analyzing individual-level activity, we find that most participants contribute only once and with little effort, leaving the top 10% of contributors responsible for almost 80% of total classifications. We also provide insights into underlying dynamics and mechanisms such as the role of effort versus speed in distinguishing top contributors, or the increase in efficiency as users participate for multiple sessions.

To the extent that the observed patterns are typical of crowd science projects more generally, they suggest that projects can indeed realize significant savings in labor cost, as well as benefits of speed due to the parallel work of a large number of individuals. However, these benefits may be easier to generate for some projects than for others. First, distributing work to the crowd is easier if projects can be decomposed into many small independent tasks such as the coding of individual pieces of data (35, 36). In contrast, complex tasks such as proving mathematical theorems require more coordination among individuals, likely limiting the number of participants who can get involved efficiently. Second, the high variability in project contributions over time suggests that crowd science may be more appropriate for tasks where the exact timing of effort is less important. For example, it is likely easier to ask the crowd to inspect a certain amount of archival data than to ask it to continuously monitor a stream of real-time data. Finally, challenges may arise for projects that require sustained participation by the same individuals. For example, efficiency in some tasks increases over time due to learning (see ref. 37 and our analyses above), but if most users only participate once, these efficiency benefits will remain unexploited. Short-term participation of individual contributors may also pose problems for projects seeking to rely on particular volunteers to collect observational data over an extended period. Of course, project organizers can try to overcome such challenges in various ways, such as through task redesign (38), technical infrastructure, or mechanisms to attract and retain users. Although a detailed consideration of these issues is beyond the scope of this article, we will highlight some opportunities below.

The observed patterns of contributions suggest that project organizers need to actively attract new users throughout the life of a project and find mechanisms to reengage existing users to reduce the rate of dropouts. Thus, although crowd science can allow substantial savings in labor costs, organizers likely require significant resources to build infrastructure that makes participation easy and fun, as well as time to engage with the user base and recruit new users (15, 29). Although these financial and nonfinancial costs are likely to vary across projects, they may be considerable. As such, scientists who consider using the crowd need to carefully compare the potential benefits and costs of crowd science with those of alternative mechanisms such as traditional laboratory structures or even paid crowd labor (e.g., AMT). Although there is currently little data on the relative efficiency of these different mechanisms, we hope that our estimates of counterfactual costs as well as of important parameters such as the rate of dropouts or the volume of effort per user in crowd science projects can help organizers make more informed decisions.

Our analysis uses data from the largest existing crowd science platform Zooniverse. This platform relies primarily on users' intrinsic motivation and interest in particular tasks or topics and clearly highlights scientific research as its main mission. There are other crowd science projects such as Foldit, EyeWire, or Phylo that rely less on users' interest in science and develop additional mechanisms to engage users, in particular by using "gamification." Among others, this may involve competition between individuals or teams of users, scoring and ranking systems, and prizes and rewards. Our results provide a useful baseline regarding contribution patterns when participation is primarily based on users' interest in the project itself, yet future

research is needed to examine contributions in projects that use additional incentives. Recent work examining gamification suggests that the qualitative patterns are likely to hold. For example, a recent case study of the game Citizen Sort shows that most players quickly lost interest in the game and that a small number of users were responsible for a disproportionate share of the total time played (29). Similarly, a very small number of users are responsible for most of the score improvements in Foldit and Phylo (7, 9). At the same time, it is likely that specific parameters such as the rate of dropouts or the exact value of Gini coefficients differ in projects using incentives or gamification, suggesting the more general need for future comparative studies using data from a broader range of projects (see also *SI Text* and *Fig. S7*).

More generally, future research is needed on the mechanisms that organizers can use to increase the effectiveness and efficiency of crowd science projects. Gamification is one such mechanism related specifically to user motivation, but research is needed on a broader range of issues such as different mechanisms to reach out to potential users (e.g., social media vs. email approaches), mechanisms to coordinate user contributions in collaborative projects, or algorithms to increase data quality while reducing the need for labor inputs (22).

A related important question is how crowd science can be leveraged for tasks that are more complex, involve interactions among contributors, or require access to physical research infrastructure. Projects such as Polymath, Galaxy Zoo Quench, Foldit, and EteRNA are pushing the boundaries in this direction (6, 7, 39). In Foldit, for example, players predict complex protein structures and can do so individually or in teams, often drawing on problem-solving approaches or partial solutions contributed by others (7). Successful attempts to crowdsourcing complex problems have also been reported outside the realm of science. For example, researchers have examined how the crowd can get involved in programming, decision making, or in the writing of nonscientific articles (28, 38, 40, 41). Many of these approaches decompose the complex task into smaller microtasks, distribute these microtasks to a large number of contributors, and then use sophisticated mechanisms to integrate the resulting contributions (28, 40). Our insights into the activities of contributors performing microtasks should be particularly relevant for such projects relying on task decomposition. Other projects offer simple tasks as a gateway for contributors to tackle more challenging assignments later in their life cycle (34, 42). Such a "career progression" requires that users remain engaged in a project over a longer period, highlighting again the need to understand return behavior and user retention.

Although some observers argue that there is plenty of "idle time" that could be used in crowd-based efforts, there may be a limit to the number of people who are willing to contribute to crowd science or to the amount of time they are willing to spend (43). As such, another important question for future research is how participation develops as an increasing number of projects compete for users' attention, including projects that address very similar scientific problems. Relatedly, research is needed on the benefits of hosting projects on larger platforms such as Zooniverse, Open-Phylo, or the Cornell Lab of Ornithology. On the one hand, projects that start on existing platforms may benefit from a large installed user base, as evidenced by the high levels of activity in Zooniverse projects in the first days of their operation. Similarly, platforms may allow organizers to share a brand name and sophisticated infrastructure that improves the user experience, potentially increasing the supply of volunteer labor while also allowing organizers to spread the fixed costs of technology over a larger number of projects. At the same time, however, users may face lower search and switching costs than for stand-alone projects, potentially exacerbating competitive effects.

Finally, to what extent may crowd science replace traditional scientific research or reduce the need for research funding? We suggest that involving the crowd may enable researchers to

pursue different kinds of research questions and approaches rather than simply replacing one type of labor (e.g., graduate students) with another (volunteers). Moreover, the interest of the crowd may be limited to certain types of projects and topics, e.g., projects that relate to existing hobbies such as astronomy and ornithology or that address widely recognized societal needs (e.g., health, environment). In those areas, however, funding agencies and policy makers should explore whether resources can be leveraged by supporting individual crowd science projects or platforms and related infrastructure.

Although our analysis is only a first step toward empirically examining individuals' participation patterns, we hope that the results contribute to a better evidence-based understanding of the potential benefits and challenges of crowd science. Such insights, in turn, are important as scientists and policy makers consider whether, how, and for what kinds of problems crowd science may be a viable and effective mode of organizing scientific research.

Materials and Methods

Detailed background information on the Zooniverse platform is available at www.zooniverse.org/about and in a number of case studies (3, 11, 23). We obtained from the Zooniverse organizers data from the seven projects that were launched in the year 2010. Table S1 summarizes key characteristics of the projects such as the project name, scientific field, type of raw data

presented to users, description of the crowdsourced activity, and the start and end dates of our observation window. For comparability across projects, we studied each project for a period of 180 d starting with the first day of its operation. We identified official launch dates based on Zooniverse blog entries and information from Zooniverse organizers.

We analyzed individual daily records for all contributors who joined one of the seven projects during the observation window. In total, there were 100,386 project participants whose activity is observed for a minimum of 1 d and a maximum of 180 d. For each contributor, we analyzed daily records of the time spent working on the project and the number of classifications made. Activity can be attributed to particular contributors because first time users create a login and password (a process that takes only a few seconds) and are asked to login before working on a project. Although some individuals may participate in multiple projects, our unit of observation is the participant in a given project, i.e., activities are recorded independently in each of the projects.

Details on the measures used and statistical analyses performed are provided in *SI Text*.

ACKNOWLEDGMENTS. We thank Chris Lintott, Stuart Lynn, Grant Miller, and Arfon Smith for providing access to data and for many valuable insights into the Zooniverse. We received extremely useful feedback from several colleagues and crowd science organizers including Raechel Bianchetti, Benjamin Mako Hill, Kourosh Mohammadmehdi, Nathan Prestopnik, Christina Raasch, Cristina Rossi-Lamastra, Paula Stephan, and Jerome Waldispuhl. Two reviewers provided exceptionally constructive comments and played an important role in shaping this paper. We gratefully acknowledge support from the Alfred P. Sloan Foundation Project "Economics of Knowledge Contribution and Distribution."

- Wiggins A, Crowston K (2011) From conservation to crowdsourcing: A typology of citizen science. *44th Hawaii International Conference on System Sciences (HICSS)*, (IEEE, New York), pp 1-10.
- Franzoni C, Sauerermann H (2014) Crowd science: The organization of scientific research in open collaborative projects. *Res Policy* 43(1):1-20.
- Nielsen M (2011) *Reinventing Discovery: The New Era of Networked Science* (Princeton Univ Press, Princeton).
- Cardamone C, et al. (2009) Galaxy Zoo green peas: Discovery of a class of compact extremely star-forming galaxies. *Mon Not R Astron Soc* 399(3):1191-1205.
- Khatib F, et al.; Foldit Contenders Group; Foldit Void Crushers Group (2011) Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nat Struct Mol Biol* 18(10):1175-1177.
- Lee J, et al. (2014) RNA design rules from a massive open laboratory. *Proc Natl Acad Sci USA* 111(6):2122-2127.
- Cooper S, et al. (2010) Predicting protein structures with a multiplayer online game. *Nature* 466(7307):756-760.
- Malone TW, Klein M (2007) Harnessing collective intelligence to address global climate change. *Innov (Camb, Mass)* 2(3):15-26.
- Kawrykow A, et al.; Phylo players (2012) Phylo: A citizen science approach for improving multiple sequence alignment. *PLoS ONE* 7(3):e31362.
- Jeppesen LB, Lakhani K (2010) Marginality and problem-solving effectiveness in broadcast search. *Organ Sci* 21(5):1016-1033.
- Christian C, Lintott C, Smith A, Fortson L, Bamford S (2012) Citizen science: Contributions to astronomy research. arXiv:1202.2577.
- Raddick MJ, et al. (2013) Galaxy zoo: Motivations of citizen scientists. *Astron Educ Rev* 12(1).
- van Knippenberg D, De Dreu CK, Homan AC (2004) Work group diversity and group performance: An integrative model and research agenda. *J Appl Psychol* 89(6):1008-1022.
- Pimm SL, et al. (2014) The biodiversity of species and their rates of extinction, distribution, and protection. *Science* 344(6187):1246-1252.
- Druschke CG, Seltzer CE (2012) Failures of engagement: Lessons learned from a citizen science pilot study. *Appl Environ Educ Commun* 11(3-4):178-188.
- Jordan RC, Ballard HL, Phillips TB (2012) Key issues and new approaches for evaluating citizen-science learning outcomes. *Front Ecol Environ* 10(6):307-309.
- Couch J (2014) Presentation given January 2014. Available at dpcpsi.nih.gov/sites/default/files/Citizen_Science_presentation_to_CoC_Jan31.pdf. Accessed May 13, 2014.
- US Government (2013) The open government partnership: Second open government national action plan for the United States of America. Available at whitehouse.gov/sites/default/files/docs/us_national_action_plan_6p.pdf. Accessed May 13, 2014.
- Google, Inc., Global Impact Awards (Zooniverse). Available at google.com/giving/global-impact-awards/zooniverse/. Accessed May 13, 2014.
- The White House, Office of Science and Technology Policy (2013) Seeking Stellar "Citizen Scientists" as White House Champions of Change. Available at whitehouse.gov/blog/2013/04/23/seeking-stellar-citizen-scientists-white-house-champions-change. Accessed May 13, 2014.
- Bonney R, et al. (2014) Citizen science. Next steps for citizen science. *Science* 343(6178):1436-1437.
- Simpson E, Roberts S, Psorakis I, Smith A (2013) Dynamic Bayesian combination of multiple imperfect classifiers. *Decision Making and Imperfection*, eds Guy TV, Karny M, Wolpert D (Springer, Berlin Heidelberg), pp 1-35.
- Fortson L, et al. (2011) Galaxy Zoo: Morphological classification and citizen science. arXiv 1104.5513.
- Ortega F, Gonzalez-Barahona JM, Robles G (2008) On the inequality of contributions to Wikipedia. *Hawaii International Conference on System Sciences, Proceedings of the 41st Annual Meeting (IEEE)*, pp 304.
- Wilkinson DM (2008) Strong regularities in online peer production. *Proceedings of the 9th ACM Conference on Electronic Commerce (ACM, New York)*, pp 302-309.
- Franke N, Hippel E (2003) Satisfying heterogeneous user needs via innovation tool-kits: The case of Apache security software. *Res Policy* 32(7):1199-1215.
- Panciera K, Halfaker A, Terveen L (2009) Wikipedians are born, not made: A study of power editors on Wikipedia. *Proceedings of the ACM 2009 International Conference on Supporting Group Work (ACM, New York)*, pp 51-60.
- Kittur A, Smus B, Khamkar S, Kraut RE (2011) Crowdforge: Crowdsourcing complex work. *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (ACM, New York)*, pp 43-52.
- Prestopnik NR, Crowston K, Wang J (2014) Exploring data quality in games with a purpose. *iConference 2014 Proceedings (iSchools)*. Available at ideals.illinois.edu/handle/2142/47311. Accessed December 19, 2014.
- Gastwirth JL (1972) The estimation of the Lorenz curve and Gini index. *Rev Econ Stat* 54(3):306-316.
- Singh PV, Fan M, Tan Y (2007) An empirical investigation of code contribution, communication participation, and release strategy in open source software development. Available at arnop.unimaas.nl/show.cgi?fid=16233. Accessed December 19, 2014.
- Ghosh RA, David P (2003) The nature and composition of the Linux kernel developer community: A dynamic analysis. *MERIT Working Paper*.
- Varshney LR (2012) Participation in crowd systems. *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference (IEEE)*, pp 996-1001.
- Jackson C, Østerlund C, Crowston K, Mugar G, Hassman KD (2014) Motivations for sustained participation in citizen science: Case studies on the role of talk. *17th ACM Conference on Computer Supported Cooperative Work & Social Computing*.
- Haessler C, Sauerermann H (2014) The anatomy of teams: Division of labor in collaborative knowledge production. Available at papers.ssrn.com/sol3/papers.cfm?abstract_id=2434327. Accessed December 19, 2014.
- Simon HA (1962) The architecture of complexity. *Proceedings of the American Philosophical Society* 106(6):467-482.
- Kim JS, et al.; EyeWires (2014) Space-time wiring specificity supports direction selectivity in the retina. *Nature* 509(7500):331-336.
- Baldwin CY, Clark KB (2006) The architecture of participation: Does code architecture mitigate free riding in the open source development model? *Manage Sci* 52(7):1116.
- Polymath DHJ (2012) A new proof of the density Hales-Jewett theorem. *Ann Math* 175(3):1283-1327.
- Bernstein MS, et al. (2010) Soylent: A word processor with a crowd inside. *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (ACM, New York)*, pp 313-322.
- Dow S, Kulkarni A, Klemmer S, Hartmann B (2012) Shepherding the crowd yields better work. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (ACM, New York)*, pp 1013-1022.
- Kwak D, et al. (2013) Open-Phylo: A customizable crowd-computing platform for multiple sequence alignment. *Genome Biol* 14(10):R116.
- Cook G (2011) How crowdsourcing is changing science. *The Boston Globe*, Nov. 11. Available at bostonglobe.com/ideas/2011/11/11/how-crowdsourcing-changing-science/dWL4DGWmq2YonHK8uOXZn/story.html. Accessed December 19, 2014.