How thanking peers sustains volunteer participation in public goods: parallel field experiments in four Wikipedia language communities

J. Nathan Matias[1,*], Julia Kamin[2], Reem Al-Kashif[1], Wanqing Psyche He[3], Max Klein[1], and Eric Pennington[4]

This manuscript was compiled on May 1, 2026

Each year, billions of people volunteer unpaid time and labor to public goods that anyone online or in their communities can benefit from. When people contribute to a common-pool resource such as the global encyclopedia Wikipedia, it is sometimes through upstream reciprocity, where people give without expecting something in return. Social scientists have hypothesized that gratitude behaviors, such as thanking others, play a role in sustaining and extending this cooperation, independently of reputation. We tested this hypothesis in a field experiment with Wikipedia communities in the Arabic, German, Polish, and Persian language editions. We randomly assigned 15,558 vetted Wikipedia contributors to receive private messages of thanks from other contributors. We then observed their continued contributions and the thanks they sent to others. Assignment to receive thanks increased the time participants contributed to Wikipedia by 11% and increased two-week contributor retention by 2.2 percentage point on average. Receiving thanks also changed how participants expressed appreciation to others. Being assigned to receive thanks also caused participants to thank others 61% more, with 99.8% of thanks sent upstream to a third party who had not previously exchanged thanks with them. Overall, this study provides experimental evidence in the field that private expressions of gratitude maintain participation in voluntary endeavors and cause recipients of gratitude to thank others upstream.

Gratitude | Public Goods | Upstream Reciprocity | Volunteering | Wikipedia | Field Experiment

Significance Statement

How can society sustain public goods that anyone can enjoy? Globally, 2.1 billion people volunteer each month (23). Volunteers sometimes contribute through "pay-it-forward" upstream reciprocity without expecting direct benefits in return. Upstream and other kinds of reciprocity have created Wikipedia, a free encyclopedia available in over 330 languages. Scientists hypothesize that gratitude is one force that could sustain upstream reciprocity. In this study, hundreds of peers sent private thank-you messages to thousands of Wikipedia contributors across four languages and two continents. A thank-you from peers increased volunteering time and volunteer retention. Thanks recipients also went on to thank others who had not previously thanked them. These findings advance the psychology of gratitude and reciprocity and offer insights to anyone organizing volunteers.

How do private expressions of thanks spread through social groups that produce public goods, and what role does that appreciation play in the formation and maintenance of those goods? Questions about the function and structure of gratitude are part of an enduring scientific discussion about the origin and maintenance of human cooperation (1–4). Pragmatically, this science seeks to understand the emergence of large-scale public goods, which are often sustained by volunteers, who contribute the equivalent of 61 million full-time workers globally each year (5). For example, tens of millions of volunteers * have created and sustained Wikipedia, a free resource of global knowledge (6–8) that has become a basic building block for education (9), elections (10), financial markets (11), and artificial intelligence (12). By investigating the role of gratitude expressions in large-scale cooperation, scientists can explain how volunteer-supported public goods come into being and also inform their effective creation and maintenance.

Social scientists describe gratitude as an emotion linked with interpersonal behaviors that form a wider structure of exchange. In this model, people receive a benefit, feel the positive affect of gratitude, and respond to that affect by expressing thanks and offering benefits themselves (3, 4, 13–15). This exchange of goods and gratitude forms the larger structures of reciprocity that constitute cooperation. In direct reciprocity (Fig 1.A), receivers express thanks to acknowledge their indebtedness to the giver and signal their willingness to return the favor in the future (3, 16). In generalized exchange (Fig 1.B, 1.C), people contribute to others who have never helped them and may never directly return the favor (4). Volunteer-produced public goods like Wikipedia are form of generalized exchange where people contribute to a common resource that others may draw from without being required to contribute (6, 17).

Social scientists have imagined different roles for gratitude in generalized exchange depending on the structure of

*https://en.wikipedia.org/wiki/Wikipedia:Wikipedians

Author affiliations: [1]Cornell University, Department of Communication, Ithaca, USA; [2]Prosocial Design Network, USA; [3]Cornell University, Department of Information Science, Ithaca, USA; [4]Harvard University, Boston, USA; *nathan.matias@cornell.edu

Writing–review & editing: all authors. Conceptualization: all except W.P.H. Methodology: all except R.A.K. Validation: all except M.K. Investigation: J.N.M., J.K., R.A.K., W.P.H. Data curation: J.N.M., W.P.H., M.K., E.P. Supervision: J.N.M., J.K., R.A.K., E.P. Software: J.N.M., M.K., E.P. Visualization: J.N.M., W.P.H., M.K. Writing–original draft: J.N.M., W.P.H. Project administration: J.N.M., J.K. Resources, funding acquisition: J.N.M. Community partnerships: J.K. Community liaison: R.A.K.

No competing interests to declare.

[2]To whom correspondence should be addressed. E-mail: nathan.matias@cornell.edu

exchange and whether appreciation is public or private. Lab and field experiments have shown that people contribute to public goods in structures of indirect reciprocity when public appreciation grows a person's reputation, increasing the chance that others will help them (Fig 1.B) (4, 18). Scholars have also studied upstream reciprocity (Fig 1.C), where people "pay it forward" to others in structures that do not loop back to benefit them, beyond the resources of the public good (15). Observational studies have found that private, upstream appreciation occurs in large-scale computer-mediated cooperation networks, and that some people will even resist the introduction of reputation mechanisms in those settings (19). Lab experiments have shown that feelings of gratitude can initiate and help maintain this private upstream reciprocity even with the reduced personal gain it entails (4, 20, 21). Yet despite evidence from simulations, surveys, and lab experiments, the causal role of gratitude and thanking behaviors in upstream reciprocity has not been validated in the field (22).

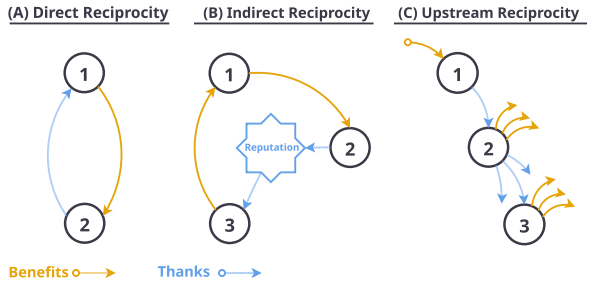

Fig. 1. Social scientists have studied the role of gratitude and thanks behaviors in (A) pairs of direct reciprocity, (B) networks of downstream (or indirect) reciprocity that come back to benefit contributors through the reputation mechanisms, and (C) networks of upstream (or general) reciprocity where benefits may not return to the giver aside from the general benefits of drawing from a public good.

How might these exchanges of private, upstream appreciation (Fig 1.C) contribute to the development and maintenance of public goods and how are they sustained? For someone to sustain repeated contributions, they need self efficacy and a belief in the utility of those contributions (4). Receiving thank-you messages can increase self-efficacy, a person's belief in their ability to accomplish a goal, by providing independent information on the quality and utility of their contributions (24, 25). This information matters in computer-mediated cooperation, where exchanges sometimes lack the social cues that otherwise provide such information (26). Expressions of gratitude may also introduce positive affects that encourage generosity (15, 27, 28). When someone feels this positive affect and receives positive feedback on the quality and utility of their contributions, they might increase the effort they contribute to cooperation and how long they continue to do so.

A full explanation of gratitude in upstream reciprocity also needs to account for structures where people thank others upstream upon receiving thanks (19). Self-efficacy about one's own contributions does not obviously explain why someone would be more motivated to evaluate and appreciate others' contributions. Instead, receiving thanks may shift how people orient to the community by elevating positive affect (emotion contagion) (29), making norms of appreciation more salient (30), changing how people see others' contributions (social appraisal) (31), or increasing attention to others' valuable work. If receiving thanks does cause people to thank others, whatever the mechanism, such an effect could further explain how gratitude behaviors both sustain contributions and extend to others in upstream reciprocity.

These questions about the effects of private gratitude expressions in upstream reciprocity have pragmatic implications for volunteer-supported public goods such as Wikipedia, a multilingual encyclopedia available in over three hundred languages.[†] Millions of Wikipedia contributors have coordinated to create one of the world's most consulted information sources,[‡] producing an estimated billions of dollars to the U.S. economy per year according to economists (32). Maintaining this free, common resource requires generalized exchange from volunteers who pay forward the knowledge they have received in the hope that others' knowledge will provide them with value in turn (6). While Wikipedia is partially maintained through the reputational benefits of downstream reciprocity (18, 33, 34), website logs also show that private thank-you messages are commonly exchanged in structures of upstream reciprocity. By mid-2018, Wikipedia contributors were sending nearly seventy thousand of these private messages per month. Consequently, Wikipedians have been as interested as social scientists to understand effects of private thank-you messages on people's upstream contributions as well as the forces that sustain a culture of appreciation.

Wikipedia is a fruitful context for field experiments on gratitude behaviors in upstream reciprocity. In the social and behavior sciences, field experiments are used to validate theories in naturalistic settings and observe unexplained causal phenomena for further theory-building (35, 36). As a cooperative endeavor sustained through generalized exchange, Wikipedia provides a naturalistic setting in which upstream, private appreciation occurs. Since the Wikipedia software hosts hundreds of distinct language encyclopedias, scientists can conduct parallel replications of a research question across multiple cultures and languages (37).

Wikipedia also includes features for sending and receiving private appreciation. When someone receives a "Thanks" on Wikipedia, they receive a private alert on the Wikipedia website (Figure 2), and an email (by default) from the sender. The message names the account that expressed thanks and identifies the specific contribution that merited thanks. Compared to public recognition on Wikipedia (18), these interpersonal messages of appreciation are private and are only easily visible by the recipient (38), limiting the role of reputation as a mechanism for their effects. The software behind Wikipedia also records the sender, receiver, time, and object of every expression of thanks sent between contributors (39).

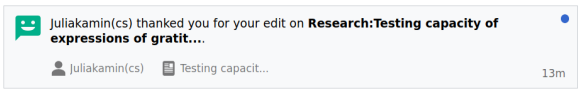

Fig. 2. Notification received by participants who received thanks

In a large-scale field experiment on Wikipedia, we tested the hypotheses that receiving thanks increases people's contributions to public goods and also causes recipients to thank

[†] https://meta.wikimedia.org/wiki/List_of_Wikipedias, accessed July 1, 2020
[‡] https://www.alexa.com/siteinfo/wikipedia.org, accessed July 1, 2020

others upstream. We designed this study in collaboration with liaisons from the Arabic, German, Persian, and Polish language communities, which collectively have over 27,000 monthly active contributors (Table 1). Working with humans and machine learning models to identify newcomers and experienced editors whose contributions merited appreciation, we randomly assigned over fifteen thousand participants to be thanked by one of 313 volunteers or to be omitted from the thanking campaign. Over the next six weeks, we studied the effect on contributions to public goods by observing how many weeks they continued to contribute to Wikipedia and how much time they spent making contributions. We also studied upstream thanks by observing the rate at which participants went on to send expressions of gratitude to others and whether those thanks were sent upstream.

Results

We evaluated the intent-to-treat (ITT) effect of receiving thanks among 15,274 participants from August 2, 2019, through March 24, 2020. We pre-registered our analysis plan before data collection, including plans for multiple comparison adjustments (40), on the Open Science Framework (OSF).[§] Among active accounts, 55% of newcomers and 46% of experienced accounts met our criteria for inclusion based on the quality of their contributions to Wikipedia (Table 1). This analysis includes all individuals randomly assigned to the treatment condition, regardless of whether they ultimately received a thank-you message from a volunteer. Among the 7,637 participants in the treatment group, volunteers chose to thank 32% of newcomers and 43% of experienced accounts (Table 1). Our manipulation check found that among 767 survey-takers, being assigned to receive thanks increased the chance of someone remembering seven weeks later that they had received thanks ($p<0.0001$).

Assignment to receive thanks increased the two-week retention of participants by 2.2 percentage point on average ($p= 0.00093$, Figure 3A, Table 2). This effect was positive in direction across all four language communities and both experience subgroups, though statistically significant only for newcomers (Tables S6, S7).

We also examined treatment effects on contributors' average daily labor hours over the next six weeks (Table 2). Being assigned to receive thanks caused participants to contribute 11% more time than the control group ($p=0.014$) (Table 2).

Receiving thanks also caused the recipients to thank others upstream. Across all four Wikipedia communities, being assigned to receive thanks caused participants to thank others 61% more than those in the control group ($p= 0.0021$, Figure 3B, Table 2). The effect was positive in direction across languages and experience subgroups, though statistically significant primarily among newcomers (Tables S6, S7). Crucially, this increase was not driven by direct reciprocity: of the 2,552 thanks sent by participants in the treatment group, 99.8% were sent to someone other than the volunteer who thanked them. Only 0.2% of thanks followed a pattern of direct reciprocity.

[§] https://osf.io/c67rg/

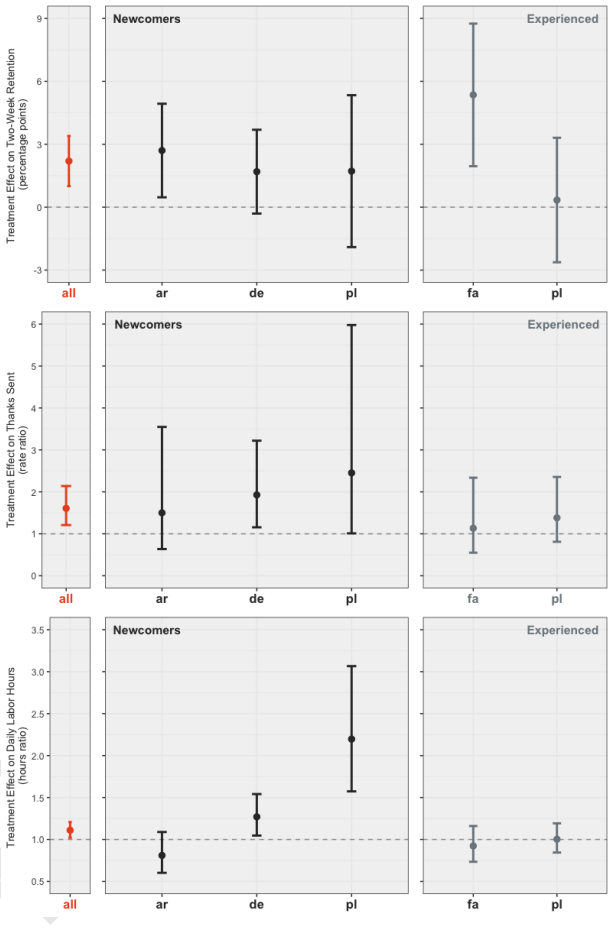

Fig. 3. ITT effects for (A) two-week retention of Wikipedia contributors, (B) thanks sent to others (incidence rate ratio, IRR), and (C) labor hours per day after treatment, in subgroups of newcomers and experienced contributors from different communities: Arabic (ar), German (de), Polish (pl), and Persian (fa) marked by ISO 639-1 language codes. Error bars represent 95% confidence intervals. Dashed lines indicate null effect (0 for percentage points, 1 for rate ratios).

Discussion

In this study, we found experimental evidence across multiple language Wikipedias for the effect of gratitude expressions to sustain, increase, and extend contributions to upstream reciprocity. In line with expectations from simulations, surveys, and lab experiments, receiving appreciation from others does cause people to "pay forward" further contributions over a longer period of time. Furthermore, this study offers novel experimental evidence that receiving private thanks causes people to express private appreciation to others in an upstream direction, largely to people who had not thanked them previously. Taken together, these findings show how private expressions of appreciation function to grow and maintain upstream reciprocity in public goods across multiple languages and cultures.

Pragmatically, this study explains one reason why Wikipedia's system for private appreciation is so widely used and what function it plays in the maintenance of this public good. Wikipedia and other largely voluntary public goods depend on many interdependent motivations to contribute. A Wikipedia article describing a work of children's fiction, for example, could receive edits from a librarian, a parent, a

Table 1. Participation rates and control group outcomes from Arabic, German, Polish and Persian Wikipedias from 2019-08-02 to 2020-02-11. "Available during the study" were accounts that were active during at least one of the waves of the study. "Included in the study" are accounts that met the quality requirements and were included in the study.

		eligible participants				average outcome values for control group			
		monthly active	available for study	included in study	% treatment thanked	two week retention rate	pre-study labor hours	daily labor hours	thanks sent
Newcomer	Polish	2414	2756	1640	14%	19%	0.027	0.012	0.038
	German	8555	10476	5434	49%	18%	0.024	0.012	0.052
	Arabic	12281	5936	3072	13%	11%	0.020	0.010	0.032
	Total	23250	19168	10146	32%	16%	0.023	0.012	0.044
Experienced	Polish	1985	8189	2996	72%	44%	0.061	0.056	0.349
	Persian	1899	4648	2416	7.3%	31%	0.028	0.035	0.264
	Total	3884	12837	5412	43%	38%	0.046	0.047	0.311

Table 2. Estimates of ITT effects among Wikipedia participants of all groups. ITT estimates were conducted with the Negative Binomial model (NegBin), the Difference in Means model (DiM), and the Tweedie model, with p-values were adjusted for multiple comparisons. For NegBin and Tweedie models, the estimate is the log-scale coefficient. The % change represents the incidence rate ratio (IRR) for NegBin and the proportional change for Tweedie. For DiM, the estimate is the difference in means.

Outcome	Estimator	N	Estimate	(% change)	Std. Error	P value
thanks sent	NegBin	15,274	0.47	(+61%)	0.14	0.0021
retention	DiM	15,274	0.022	(+2.2%)	0.0061	0.00093
labor hours per day	Tweedie	15,274	0.10	(+11%)	0.043	0.014
manipulation	DiM	767	0.29	(+29%)	0.035	<0.0001

publisher, a competing author, a young reader, or an administrator of that language edition. Prior field experiments have shown the importance of reputation and indirect reciprocity to some Wikipedia contributors, finding that public rewards can increase participation (18, 33). This study shows that upstream reciprocity plays a complementary role in growing and sustaining participation for the tens of thousands of people who send and receive private thanks each month. This finding could help Wikipedia language communities retain thousands of contributors a year, elicit many more hours of contributions from each other, and further extend the practice of appreciation across the movement.

This study also experimentally confirms what scientists have observed elsewhere, that private messages of thanks to third parties upstream are more common than reciprocal thanks (19). Prior research has described gratitude as a "moral barometer" that follows receipt of benefits, signaling a sense of obligation to the giver (3). This could still be true for public goods and systems of indirect reciprocity, where people routinely receive benefits from people with whom they have not previously helped directly. Under such circumstances, appreciation would follow the indirect or upstream structure of exchange.

As an experiment that does not study multiple rounds of thanks-sending, this study cannot support claims about contagion or cascades, nor can it document how common the practice currently is. Yet we show that receiving thanks can itself cause people to extend thanks to third parties for the benefits of upstream reciprocity they receive.

Since only 32% of newcomers and 43% of experienced accounts assigned to the treatment group received a thanks message from a volunteer, our ITT estimates constitute a lower bound on average treatment effects that are likely larger. By giving volunteers a choice, we traded precision in estimating the effect of a single thanks message for the external validity of studying contributions that volunteers considered worth thanking. Since our labor hours measure focused on contributions to pages and excluded time spent thanking others, we may have observed a larger effect on labor hours with a more inclusive measure. The study could not prevent participants in the treatment group from subsequently thanking other participants in the treatment or control groups, something we observed in some cases (see Supporting Information, Spillover Sensitivity Analysis). As is common in multi-site field experiments, we observed variation in coefficients and confidence intervals both between and within languages. While the number of groups is not large enough to systematically study that variation, possible causes of variation could include differences in inclusion criteria, community size, measurement, and spillover.

This experiment is not designed to distinguish between possible psychological mechanisms for the effects we observe. It is possible that receiving thanks grew contributions through an increase in self-efficacy (24) or a positive affect (15, 27, 28). Upstream thanking may have been influenced by such affects. By generating a message to the thanks recipients, the intervention may also have communicated a community norm of thanking that participants, especially newcomers, may not have known about previously (30). In a follow-up, exploratory analysis of labor hours, we did observe interaction effects for newcomer accounts compared to more experienced contributors, effects that are consistent with theories of self-efficacy and social norms (see Supporting Information, Exploratory Subgroup Analysis). Further field and lab research could explore those mechanisms with greater fidelity (35).

Materials and Methods

We designed a field experiment to test the influence of receiving private expressions of gratitude on pro-social behavior among contributors in Arabic, German, Persian, and Polish Wikipedias. We tested whether receiving thanks could cause volunteer contributors of different experience levels to increase participation and send appreciation to others in turn. The relevant governance entities for each language Wikipedia consented to participate (see Supporting Information, Community Consent and Governance). Participants were debriefed after the study, in a protocol approved by Princeton University (11587) and considered exempt by the Cornell University Institutional Review Board (1912009282). Full study materials are available in the Supplementary Materials and OSF.[¶]

Within language Wikipedia groups, we used machine learning and human-coded criteria to identify Wikipedia editors whose contributions were eligible for appreciation. Communities considered participants eligible for the study if they made at least four eligible contributions to Arabic, German, or Polish Wikipedia in the previous 90 days. Persian Wikipedians selected a window of eligibility of 540 days. Edits were eligible if they were made to pages other than the account's personal profile, were not discussion comments, and met language-specific criteria. In German, edits were eligible if they were a "flagged revision", an edit that has been marked useful by a separate expert Wikipedian. In Arabic, Persian, and Polish Wikipedias, we used the ORES machine learning system, which had been trained by community members to identify good-faith and non-damaging edits (Table S3) (41). Some communities limited their efforts to newcomers or experienced accounts only, to ensure their volunteers could thank a sufficient number of participants. Among active accounts, 55% of newcomers and 46% of experienced accounts met the criteria and were included (Table 1; see Sampling Criteria for full details).

Participants were stratified by experience level (newcomers who had joined in the previous 90 days versus experienced editors with more experience) and by language (Arabic, German, Persian, and Polish). Within each experience–language stratum, we randomly assigned participants to receive a thanks message or serve as a control using a matched-pairs design. Pairs were formed by matching participants on the number of thanks messages they had sent previously and their labor hours contributed over the previous 84 days. Within each matched pair, one unit was randomly selected to receive treatment and the other to serve as the control. This design minimizes the variance of the difference-in-means estimator of the average treatment effect (42). Before analysis, we excluded matched pairs in which the treated member did not receive exactly one thanks message due to software errors, account deletion, or volunteer decisions to skip a participant. These are protocol-level exclusions of matched pairs, not post-hoc removal of individual outcomes. Table S5 reports pre-dropped results with these pairs.

We created a software application to encourage experienced Wikipedia volunteers to send a single message of appreciation to participants in the treatment group, thanking them for a specific contribution (Figure S2). When an experienced Wikipedian using the software chose to send thanks, a message would be sent immediately to the recipient *on behalf* of the sender through the Wikipedia interface (Figure 2). The message would also arrive via email if the recipient allowed email notifications (the default on Wikipedia).

From August 2, 2019, through February 11, 2020, a total of 313 volunteers reviewed a queue of eligible edits by participants in the treatment group. To meet our sample size goals for newcomers, we conducted the study in a series of 84-day waves, including in each wave accounts created in the period since the previous wave. After removing 142 matched pairs associated with software errors and opt-out requests as described above, the study included a total of 15,274 participants.

We then observed participant behavior on Wikipedia over the next six weeks after they (or their matched pair) received thanks, up through March 24, 2020. In the case of pairs that did not receive thanks, we observed the six weeks after random assignment. Using public data from the Wikipedia software system, we created measures of two-week retention, the difference in labor hours they contributed, and the number of thanks they sent to others during the observation period. After data collection was complete, we debriefed all participants in the study and granted them the opportunity to opt out of inclusion in the data analysis if they wished.

The number of thanks sent over six weeks is a count variable drawn from Wikipedia's public records for a given participant. We estimate the average treatment effect on the incidence rate of thanks sent using a negative binomial regression model. We report the exponentiated coefficient as an incidence rate ratio (IRR) (43), and express it as a percentage change, $(IRR - 1) \times 100$.

Two-week retention is a binary variable that records whether a participant made at least one edit during the post-treatment observation period, measured from the second week onwards. This measure is designed to infer whether someone is still actively contributing to Wikipedia at that time or is no longer active. The measure omits observations in the period before the second week started because we wish to know if the contributor was still active in the period after the first week of observation concluded. The measure includes a window of observation beyond the end of the week - if someone comes back to edit later, we consider them still active in their second week. In this case, the measurement period starts at the beginning of the second week after random assignment and extends five more weeks to identify contributors who have discontinued editing, for a total observation period of 42 days. If a participant makes no contributions in the five-week period, they are classified as having stopped contributing. We record two-week retention using this five-week window because some contributors to Wikipedia only make edits every few weeks. Without such a window, the outcome variable would undercount contributors who make edits less frequently than every week.

We measure the duration of contributions to Wikipedia by inferring labor hours, an aggregate estimate of the time spent contributing to Wikipedia across edit sessions. These sessions are inferred from contiguous sequences of contributions near to each other in time, with a minimum of 430 seconds for a

[¶] https://osf.io/j64vs/

single-edit session (44). This measure does not include time spent thanking others.

As a manipulation check, we also surveyed participants after six weeks to ask if they remembered receiving thanks. Among participants, 5% completed the survey, which was balanced between treatment and control. We pre-registered hypotheses based on other survey responses, but withdrew those questions upon observing the low response rate.

This study pre-registered an ITT estimate using a difference-in-means estimator for two-week retention and the manipulation check that accounted for randomization blocks (45). These ITT estimates permit organizations to evaluate the pragmatic outcome of organizing a Wikipedia-wide intervention, accounting for selection by thankers. To estimate the effect on the number of thanks sent, we pre-registered an ITT negative binomial estimator that adjusted for language and newcomer status. This model includes random intercepts for the matched pair design (Table S1). ‖

The pre-registered analysis for the labor hours outcome estimated the average treatment effect on the difference in labor hours before and after treatment. We later identified two issues with this specification. First, this difference-in-difference analysis obscured the distinction between high-duration contributors and those who participate at levels close to the minimum observable time. Second, an analysis of residuals in this difference-in-means model revealed that the specification did not fit the zero-inflated and highly skewed data well (Figure S5). To address this issue, we adjusted the outcome variable and the model. First, we used a measure of post-treatment labor hours per day as the outcome variable. Second, to address the skewed nature of labor hours per day, we fit generalized linear mixed models (GLMM) with the Tweedie distribution (46) (Figure S6). Full subgroup-specific interpretations of these effects are provided in the Supporting Information (Exploratory Subgroup Analysis; Table S8, Figure S3).

The pre-registration also included an analysis to estimate the Complier Average Causal Effect (CACE). We decided to drop this analysis due to the potential selection bias introduced by preferences among volunteers for the kinds of edits they chose to thank.

Sampling Criteria. We identified potential thanks recipients through a two-stage process that started with automated systems to create a list of eligible candidates who were randomized into the study. Treatment group candidates were then shown to volunteer thankers, who could choose whether to send appreciation to a given study participant. All participants that were assigned randomization were included in the study, whether or not thankers chose to skip them, a choice that sets up the ITT study design. We conducted multiple waves of assignment. If not enough participants were available in the first wave, we waited for further newcomers to join and repeated the above procedure. Inclusion criteria are described in detail in the pre-analysis plan. ‖

Based on a simulation-based power analysis, we estimated the required sample sizes to achieve an 80% probability of detecting a statistically significant effect for each dependent variable. The minimum detectable effects used in the power analysis were: an additive increase of 0.25 hours (15 minutes) in labor hours over 90 days, a 25% relative increase in the probability of remaining active at two and four weeks, and an average of 0.1 additional thanks sent per participant over 90 days. We set the goal of including 1,750 newcomers in Arabic, 3,000 newcomers in German, and 800 newcomers in the Polish-language Wikipedia. Among experienced accounts, we estimated that 2,400 would be needed on the Polish-language Wikipedia and 2,400 on the Persian-language Wikipedia.

Sampling for this study was conducted on an ongoing basis throughout the study period. Accounts were eligible for inclusion in the study if they were considered active: they must have contributed to Wikipedia in the previous 90 days and have made at least four high-quality contributions to a given language Wikipedia in the previous 540 days. These criteria excluded previous thanks sent and user-page edits, which are a person's contributions to their own profile on Wikipedia. For example, if three of an account's four contributions were to their own profile in the previous 540 days, they would not be included in the study.

For Persian, Polish, and Arabic language Wikipedias, our inclusion criteria included judgments of "good faith" edits by the ORES system, a machine learning model trained by Wikipedia community members to classify contributions to Wikipedia. At the time of this study, this "algorithmic scoring service," consisted of 102 machine learning classifiers operating in 41 languages that had been trained by volunteers in each language group. We made use of a "good faith" edit classifier that could review an individual's contribution to Wikipedia and provide a probability on a scale of 0 to 1 that the edit would be judged as good faith by human reviewers. Evaluations of these models by the Wikimedia Foundation show that they exhibited a precision of at least 0.98 and a recall of at least 0.99 across all models (Table S3).

Participants in the study were labeled as newcomers if they had created the account within the previous 90 days and were labeled as experienced if they had created the account more than 90 days before.

Recruiting Thankers. We recruited a total of 313 thanks senders through a series of banner ads and surveys that were placed at the top of German, Polish, Arabic, and Persian language Wikipedias starting June 9, 2019. These banners were visible to logged-in Wikipedia accounts that were browsing the site and that met the inclusion criteria. German and Polish Wikipedia accounts were recruited if they had been granted permission by the community to "flag revisions," giving them the power to choose which version of an article was displayed to readers by default. Persian Wikipedia accounts were included if they were registered for at least a year and had contributed at least 500 edits. Arabic Wikipedia contributors were shown the banner if the community had granted them "autoreview" permissions, indicating that they have earned enough trust that their contributions would not be flagged for review by other editors. Per-community recruitment capacity, the resulting volunteer-thanking constraints, and the language-specific subgroup allocation (e.g. no experienced cohort in Arabic or German Wikipedia) are documented in the Supporting Information (Cross-Language Implementation).

Selecting and Allocating Thanks. Thankers participated by logging into a software interface for sending thanks. This software presented them with accounts to potentially thank,

‖ https://osf.io/c67rg/

showing them the contributions that made the account eligible to be considered for thanks. The software showed the before-and-after view of four contributions made by the account and presented thankers with the option to send thanks for one of these contributions or pass over the account (Figure S2). A full description of the user interface is included on OSF. **

Receiving Thanks. Many participants who were sent a thanks message received an email from the Wikipedia platform notifying them that they had been thanked. Any participants who had previously navigated through the Wikipedia platform's settings to manually disable email notifications would have received a notification of the thanks message through a bright red indicator at the top of the page when logged in to Wikipedia (Figure 2). Participants were not aware that they were being observed for the gratitude intervention.

GDPR Compliance. All participants were sent a debriefing message after the study concluded. Four people requested their data be removed from the study, a request we complied with.

ACKNOWLEDGMENTS. This study protocol was approved by the Princeton University (11587) and exempted by Cornell University (1912009282) Institutional Review boards. This research was funded by the Templeton World Charity Foundation grant #0411. At CAT Lab, communities are co-partners in our research. We are deeply grateful to the many people in the Wikipedia community who advised and supported this work. We are especially indebted to our liaisons from partnering Wikipedias – Reem Al-Kashif, Christine Domgörgen, Mohamed ElGohary, Maria Heuschkel, Amir Ladsgroup, Wojciech Pędzich, Mohsen Salek, and Natalia Szafran-Kozakowska. These liaisons worked closely with us to design the study and guide us from start to finish. This work is a product of their endless wisdom, patience and dedication. We also appreciate the 313 volunteers who generously used our app to send thanks as part of our study. Many also helped us test and fix software bugs to improve the quality of this study. We appreciate in-kind support from Wikimedia Deutschland, who provided meeting space for community conversations. We are also grateful to the Drew Margolin, David Rand, the Lazer Lab, and the attendees of the annual Wikimania conference for feedback on early drafts of this paper, and to Edward Platt for feedback on data collection code.

1. A Smith, *The Theory of Moral Sentiments.* (J. Beatty and C. Jackson, Ireland), (1777).
2. G Simmel, Faithfulness and Gratitude in *The Gift: An Interdisciplinary Perspective.* (Amsterdam University Press), (1996) Google-Books-ID: o7wq4oO2B0YC.
3. ME McCullough, SD Kilpatrick, RA Emmons, DB Larson, Is gratitude a moral affect? *Psychol. bulletin* **127**, 249 (2001) Publisher: American Psychological Association.
4. B Simpson, A Harrell, D Melamed, N Heiserman, DV Negraia, The Roots of Reciprocity: Gratitude and Reputation in Generalized Exchange Systems. *Am. Sociol. Rev.* **83**, 88–110 (2018).
5. UN Volunteers, What is not counted does not count: Global volunteering estimates in *State of the World's Volunteerism Report 2022: Building Equal and Inclusive Societies.* (United Nations), (2021).
6. Y Benkler, *The wealth of networks: How social production transforms markets and freedom.* (Yale University Press), (2006).
7. J Reagle, J Koerner, eds., *Wikipedia @ 20: Stories of an Incomplete Revolution.* (The MIT Press), (2020) Accepted: 2020-12-15T13:31:29Z.
8. D Jemielniak, *Common Knowledge?: An Ethnography of Wikipedia.* (Stanford University Press), (2020).
9. C Okoli, M Mehdi, M Mesgari, Fielsen, A Lanamäki, Wikipedia in the eyes of its beholders: A systematic review of scholarly research on Wikipedia readers and readership. *J. Assoc. for Inf. Sci. Technol.* **65**, 2381–2403 (2014) _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.23162.
10. H Salem, F Stephany, Wikipedia: a challenger's best friend? Utilizing information-seeking behaviour patterns to predict US congressional elections. *Information, Commun. & Soc.* **26**, 174–200 (2023) Publisher: Routledge _eprint: https://doi.org/10.1080/1369118X.2021.1942953.
11. M Hinnosaar, T Hinnosaar, M Kummer, O Slivko, Wikipedia matters. *J. Econ. & Manag. Strateg.* **32**, 657–669 (2023).
12. S Baack, A Critical Analysis of the Largest Source for Generative AI Training Data: Common Crawl in *The 2024 ACM Conference on Fairness, Accountability, and Transparency.* (ACM, Rio de Janeiro Brazil), pp. 2199–2208 (2024).
13. MY Bartlett, D DeSteno, Gratitude and prosocial behavior: Helping when it costs you. *Psychol. science* **17**, 319–325 (2006) Publisher: SAGE Publications Sage CA: Los Angeles, CA.
14. Y Horita, M Takezawa, T Kinjo, Y Nakawake, N Masuda, Transient nature of cooperation by pay-it-forward reciprocity. *Sci. reports* **6**, 19471 (2016).
15. MA Nowak, S Roch, Upstream reciprocity and the evolution of gratitude. *Proc. royal society B: Biol. Sci.* **274**, 605–610 (2007) Publisher: The Royal Society London.
16. SB Algoe, Find, Remind, and Bind: The Functions of Gratitude in Everyday Relationships. *Soc. Pers. Psychol. Compass* **6**, 455–469 (2012).
17. J Fulk, AJ Flanagin, ME Kalman, PR Monge, T Ryan, Connective and communal public goods in interactive communication systems. *Commun. Theory* **6**, 60–87 (1996) Publisher: Wiley Online Library.
18. J Gallus, Fostering public good contributions with symbolic awards: A large-scale natural field experiment at Wikipedia. *Manag. Sci.* **63**, 3999–4015 (2016).
19. ES Spiro, JN Matias, A Monroy-Hernández, Networks of gratitude: Structures of thanks and user expectations in workplace appreciation systems in *Tenth International AAAI Conference on Web and Social Media.* (2016).
20. ME McCullough, MB Kimeldorf, AD Cohen, An Adaptation for Altruism: The Social Causes, Social Effects, and Social Evolution of Gratitude. *Curr. Dir. Psychol. Sci.* **17**, 281–285 (2008).
21. D DeSteno, MY Bartlett, J Baumann, LA Williams, L Dickens, Gratitude as moral sentiment: emotion-guided cooperation in economic exchange. *Emotion* **10**, 289 (2010).
22. LK Ma, RJ Tunney, E Ferguson, Does gratitude enhance prosociality?: A meta-analytic review. *Psychol. bulletin* **143**, 601 (2017).
23. 2026 State of the World's Volunteerism Report. Volunteerism and its Measurements, (United Nationas Volunteers (UNV), Bonn), Technical report (2025).

** https://osf.io/j64vs

24. A Bandura, Self-efficacy: toward a unifying theory of behavioral change. *Psychol. review* **84**, 191 (1977).
25. MK Thomsen, Citizen Coproduction: The Influence of Self-Efficacy Perception and Knowledge of How to Coproduce. *The Am. Rev. Public Adm.* **47**, 340–353 (2017).
26. JB Walther, Social Information Processing Theory (CMC) in *The International Encyclopedia of Interpersonal Communication*. (John Wiley & Sons, Ltd), pp. 1–13 (2015) _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118540190.wbeic192.
27. K Gray, AF Ward, MI Norton, Paying it forward: generalized reciprocity and the limits of generosity. *J. experimental psychology: Gen.* **143**, 247 (2014).
28. DT Kong, LY Belkin, Because I want to share, not because I should: Prosocial implications of gratitude expression in repeated zero-sum resource allocation exchanges. *Motiv. Emot.* **43**, 824–843 (2019) Publisher: Springer.
29. JH Fowler, NA Christakis, Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the framingham heart study. *Bmj* **337** (2008).
30. ME Tankard, EL Paluck, Norm perception as a vehicle for social change. *Soc. Issues Policy Rev.* **10**, 181–211 (2016).
31. B Parkinson, Interpersonal Emotion Transfer: Contagion and Social Appraisal. *Soc. Pers. Psychol. Compass* **5**, 428–439 (2011).
32. E Brynjolfsson, A Collis, How should we measure the digital economy. *Harv. business review* **97**, 140–148 (2019).
33. MA Restivo, Why wiki works: Peer production and making knowledge the wiki way. (2014).
34. T Kriplean, I Beschastnikh, DW McDonald, Articulations of wikiwork: uncovering valued work in wikipedia through barnstars in *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. pp. 47–56 (2008).
35. EL Paluck, RB Cialdini, Field research methods. *Handb. research methods social personality psychology* pp. 81–97 (2014).
36. RB Cialdini, Full-cycle social psychology. *Appl. social psychology annual* (1980) Publisher: Sage Publications.
37. BM Hill, A Shaw, Studying populations of online communities. *The Handb. Networked Commun. Oxf. Univ. Press. New York, NY* (2017).
38. JN Matias, J Kamin, M Klein, Kittens, Baklava, and Bubble Tea: How Wikipedians Thank Each Other in Different Languages (2019) Library Catalog: community.globalvoices.org.
39. S Goel, A Anderson, L Zia, Thanks for Stopping By: A Study of "Thanks" Usage on Wikimedia in *Companion Proceedings of The 2019 World Wide Web Conference*. pp. 1208–1211 (2019).
40. S Holm, A simple sequentially rejective multiple test procedure. *Scand. journal statistics* pp. 65–70 (1979) Publisher: JSTOR.
41. A Halfaker, J Morgan, A Sarabadani, A Wight, ORES: Facilitating re-mediation of Wikipedia's socio-technical problems, (Working Paper. Wikimedia Research. https://meta. wikimedia. org/wiki . . .), Technical report (2016).
42. Y Bai, Optimality of matched-pair designs in randomized controlled trials. *Am. Econ. Rev.* **112**, 3911–40 (2022).
43. JS Long, J Freese, *Regression Models for Categorical Dependent Variables Using Stata, Third Edition*. (Stata Press, College Station, TX), 3 edition edition, (2014).
44. RS Geiger, A Halfaker, Using edit sessions to measure participation in Wikipedia in *Proceedings of the 2013 conference on Computer supported cooperative work*. pp. 861–870 (2013).
45. G Blair, et al., Package 'estimatr'. *Stat* **7**, 295–318 (2018).
46. ME Brooks, et al., glmmtmb balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *The R J.* **9**, 378–400 (2017) https://doi.org/10.32614/RJ-2017-066.