Cooperation in a Peer Production Economy
Experimental Evidence from Wikipedia*

Yann Algan† Yochai Benkler‡ Mayo Fuster Morell§ Jérôme Hergueux¶

January 2014

Abstract

The impressive success of peer production – a large-scale collaborative model of production primarily based on voluntary contributions – is difficult to explain by relying solely on standard assumptions about individual preferences. This paper studies the prosocial foundations of cooperation in Wikipedia, a peer production economy in which monetary incentives play no role in shaping individual behavior. We design an online experiment coupled with observational data to elicit social motives within a representative sample of 850 Wikipedia contributors, and use those measures to predict subjects' field contributions to the Wikipedia project. We thus provide the first comprehensive field test of existing economic theories of prosocial motives for contributing to real-world public goods. We find that regular editors' field contributions to Wikipedia are strongly related to their level of reciprocity in a conditional Public Goods game and in a Trust game and to their revealed preference for social image within the Wikipedia community, but not to their level of altruism either in a standard or in a directed Dictator game. The extent of participation within the group of Wikipedia administrators – who self-selected into performing a policing role within the Wikipedia community – is positively related to their revealed preference for social image but, unlike regular contributors, strongly *negatively* related to their level of reciprocity. Using our measure of trust in strangers from the Trust game, we show that while trust is unrelated to contribution levels among regular editors, less trusting Wikipedia administrators are significantly more active and more likely to exercise their policing rights.

JEL classification: H41, C93, D01, Z13
Keywords: Field Experiment, Public Goods, Social Preferences, Peer Production, Internet

* We gratefully acknowledge financial support from the European Research Council (ERC Starting Grant) and logistical support from the Sciences Po médialab and the Wikimedia Foundation. We are grateful to Anne l'Hôte and Romain Guillebert for outsdanding research assistance. We are greatly indebted to Dario Taraborelli for his instrumental help in moving this project forward, Benjamin Mako Hill and Aaron Shaw for generously sharing some of their data on and Timo Tijhof for coding the recruitment banner in coordination with us. This paper has benefited from invaluable conversations with Nava Ashraf, Shane Greenstein, Guillaume Fréchette, Stuart Geiger, Ed Glaeser, David Laibson, Sendhil Mullainathan, Klaus Schmidt, Andrei Shleifer and Gabriel Smagghue. We thank the Berkman Cooperation group and numerous seminar participants for helpful comments. Our last thank goes to the Wikipedia community and, in particular, to the 850 Wikipedia contributors who took part in this study.

† Sciences Po, Department of Economics and CEPR. e-mail: yann.algan@sciences-po.fr

‡ Harvard Law School and Berkman Center for Internet & Society at Harvard University. e-mail: ybenkler@law.harvard.edu

§ Autonomous University of Barcelona, Institute of Government and Public Policies. Postdoctoral fellow, Berkman Center for Internet & Society at Harvard University. e-mail: mayo.fuster@eui.eu

¶ University of Strasbourg, Institute of Political Studies and Sciences Po, Department of Economics. Research fellow, Berkman Center for Internet & Society at Harvard University. e-mail: jhergueux@cyber.law.harvard.edu

"The problem with Wikipedia is that it only works in practice. In theory, it can never work."
Kizor, Wikipedia administrator.[1]

1 Introduction

Peer production is characterized by the development of large-scale, collaborative and primarily voluntary based models of production in some of the most innovative and competitive sectors of information and technology (Benkler 2002; 2006; 2013). One flagship of this "New Economy" is the impressive growth of Internet mediated cooperation for the provision of public goods. Over the past 20 years, online communities of volunteers have proven surprisingly successful at developing and freely releasing pieces of computer software and information goods of substantial economic value. Those peer produced goods increasingly compete with their firm-based and for profit counterparts.[2] Among the distinctive features of this emerging production model is the fact that individuals voluntarily self-assign and successfully coordinate work in the absence of price signals, and without any pre-specified design rule or formal leadership. Those organizational features of peer production make its success difficult to explain by relying solely on standard assumptions about individual preferences.

This paper builds upon the theory of voluntary cooperation in public goods like environments to study the prosocial foundations of cooperation in Wikipedia. We elicit the prosocial preferences of Wikipedia contributors with an online experiment coupled with observational data, and seek to relate those preferences to subjects' field contributions to the Wikipedia project. Wikipedia is a paradigmatic example of peer production and a particularly clean study site, as it is difficult to derive monetary rewards from one's contributions to the project. Unlike, e.g., open source software, this feature of Wikipedia allows us to separate out extrinsic from intrinsic motivations to contribute, and study purely the prosocial motivations aspect of peer production.[3] Besides, it is possible to reliably extract from the Wikipedia website a complete record of editors' contributions to this real-world public good, so that we can base our study on experimental and observational data rather than self-reporting.

Since its inception in 2001, Wikipedia has grown to host over 25 million freely usable articles in 285 languages. Its revealed informational value seems to be enormous to society, as it receives over 80 million unique visitors per month in the United States alone,[4] and that 60% of European doctors declare using it for professional purposes.[5] As a matter of fact, an early evaluation of the quality of Wikipedia's scientific entries found them to be practically indistinguishable from those in encyclopedia Britannica

[1] See http://www.nytimes.com/2007/04/23/technology/23link.html?ei=5124&en=435e5b69b6b3ceac&ex=13&_r=0, accessed February 2013.

[2] To name a few telling examples, the open-source web browser Mozilla Firefox is currently used by 25% of all Internet users, the open-source web server Apache serves 63% of all Internet websites, Wikipedia.org is the 5th most visited website on the Internet and the user-generated game Counter-Strike is one of the most popular and long-lasting video games of a 25 billion dollars industry in the U.S. alone. Most recently, Google's decision to release the source code of Android under an open source license so that it could be peer produced significantly accelerated its development and allowed it to catch up and overtake Apple's iOS as the dominant smartphone operating system.

[3] Empirically speaking, it is a well known fact that about 50% of open source software contributors derive some kind of monetary benefit from their contributions (see, e.g., Lerner and Tirole (2002)).

[4] See http://www.comscore.com/Insights/Press_Releases/2012/9/comScore_Media_Metrix_Ranks_Top_50_US_Web_Properties_for_August_201, accessed February 2013.

[5] See http://www.pmlive.com/find_an_article/allarticles/categories/General/2011/june_2011/features/dr_wikipedia_will_see_you_now..._280528

(Giles 2005). Despite evidence of substantial economic value, however, this peer production economy has been vastly overlooked in the economics literature so far.[6]

Every Wikipedia reader holding some private information of potential value to the encyclopedia faces a standard public goods dilemma. While it is individually costly put one's knowledge in convenient shape for the general public to use, the content contributed by others is immediately available for anyone to see and use at no cost. According to the standard rational actor model, this should lead to no contributions being made in the first place. Importantly, the cost of contributing valuable information to Wikipedia in terms of effort and time is of a different nature – and arguably higher – than the cost of contributing to, e.g., a personal blog. As nicely stated in the Wikipedia Neutral Point of View policy, "articles must not *take* sides, but should *explain* the sides, fairly and without bias. This applies to both what you say and how you say it."[7] Obviously, Wikipedia would not be considered such a useful informational resource if it was merely a place for individuals to push their own personal views. Contributors are therefore expected to communicate knowledge in an encyclopedic format, provide reliable secondary sources for their claims, and resolve disputes through constructive discussions and consensus building. Since any contributor can easily revert the contributions of any other, this laudable goal would probably go unheeded without some shared cooperative norms and prosocial standards.

Individuals' intrinsic motivations for contributing to a public good can be manifold. Economic theory, however, has mainly focused on the *prosocial* foundations of cooperation in public goods like environments. Specifically, three classes of social motives have been put forward in the theoretical literature: (i) altruistic motives, either in the form of "pure altruism" or "warm-glow" (Andreoni 1989; Andreoni 1990; Anderson et al. 1998) (ii) reciprocity motives (Rabin 1993; Dufwenberg & Kirchsteiger 2004; Falk & Fischbacher 2006) and (iii) social image motives (Holländer 1990; Bénabou & Tirole 2006; Andreoni & Bernheim 2009; Ellingsen & Johannesson 2008, 2011).

This paper is the first to comprehensively test for the relative role of each class of social motive for incentivizing sustained contributions to a real-world public good.[8] Because Wikipedia in itself works as a repeated public goods experiment, we think of it as an ideally suited field for testing the external validity of those theories. On the methodology side, this paper illustrates the potential usefulness of coupling experimental methods with computational social science techniques in order to relieve the tension between internal and external validity in economic experiments. Indeed, while it is possible to leverage large samples and achieve high internal validity with online experiments (Hergueux & Jacquemet 2013) the Internet also provides a wealth of externally valid observational data on individuals' field behavior (Lazer et al. 2009).

Based on a representative sample of 850 Wikipedia contributors, we find that measures of reciprocity and social image motives – but not altruism – are significantly associated with the trajectory of Wikipedia users from a non-contributor to an engaged contributor. Our field experiment thus shows that reciprocity and social image are both strong motives for sustaining cooperation in peer production environments, while altruism may not be. In this process, reciprocity and social image appear as

[6] One notable exception for our purpose is Zhang & Zhu (2011).

[7] See https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view, accessed December 2013.

[8] An extensive literature has investigated the role of those three classes of social motives in people's (lack of) willingness to sustain cooperation in repeated public goods games in the lab, with unequal success. There is some evidence supporting the altruistic motive, although its effect appears to be inconsistent and not quantitatively large (Andreoni 1995; Palfrey & Prisbrey 1997; Goeree et al. 2002; Andreoni & Miller 2002; Vesterlund et al. 2009). By contrast, lab experiments have provided strong evidence in support of the reciprocity motive (Burlando & Guala 2005; Gächter & Thöni 2005; Page et al. 2005; Cinyabuguma et al. 2005; Charness & Yang 2007; Gunnthorsdottir et al. 2007; de Oliveira et al. 2009b; Fischbacher & Gächter 2010). The social image motive is also supported by rather strong experimental evidence (Andreoni & Petrie 2004; Rege & Telle 2004; Ariely et al. 2009) and its role has recently been confirmed in careful field experiments (Andreoni et al. 2011; DellaVigna et al. 2012).

substitutable motivational drivers rather than complementary ones. Because social image motives are difficult to measure experimentally, we exploit the observational data that is available from the Wikipedia website to construct measures of revealed preference for social image within the Wikipedia community. Controlling for a vector of demographic variables, our estimates indicate that moving from no reciprocity to full reciprocity in a conditional Public Goods game and in a Trust game is associated with a 122% and a 211% increase in the number of Wikipedia contributions, respectively, while revealing a preference for social image is associated with a fivefold increase in the number of contributions made to the project.

Interestingly, however, our experimental measures of taste for reciprocity do not predict anymore the contribution patterns of those editors who are already "super contributors" to the Wikipedia project (i.e. those who typically exhibit more than 2,000 and up to several hundreds of thousands of Wikipedia contributions) while a taste for social image continues to do so. Within this highly engaged group, revealing a taste for social image is associated with a 30 to 33% increase in the number of contributions made to the project.

Finally, we study the contribution patterns of Wikipedia administrators. Those contributors form a conceptually distinct class of highly engaged Wikipedia contributors who opted-in a very competitive peer-review process at the end of which they were granted with special oversight rights over the encyclopedia. We find evidence that administrators who participate relatively more generally exhibit a higher taste for social image, but also a significantly *lower* taste for reciprocity. This negative correlation between prosociality levels and Wikipedia participation can be explained by the very role that those highly engaged contributors self-selected into performing. As system operators, the goal of those contributors is to keep Wikipedia up and running, which involves dealing with a high number of potentially malicious users. In this sense, we posit that some Wikipedia administrators may feel responsibility towards the system, and not the people who contribute to it. We test this hypothesis directly by using an experimental measure of general trust derived from the Trust game. We find that moving from full trust to no trust in anonymous strangers is significantly associated with a 107% increase in Wikipedia activity for this group, while it has no predictive power over the contribution patterns of regular contributors. Again, our experimental measures of general and directed altruism do not seem to predict contributions patterns within this or any other group.

This paper is related to a burgeoning stream of the literature that has begun to explore the predictive power of experimental measures of social motives on field outcomes. In his seminal work, Karlan (2005) uses the Trust game to obtain individual measures of reciprocity and shows that those can be used to predict loan repayment among participants in a microcredit program. Laury and Taylor (2008) and De Oliveira et al. (2009a) relate the propensity of their subjects to cooperate in a Public Goods game in the lab to their propensity to contribute to a charitable cause in the field. One prominent limitation of those studies, however, is that they both obtain information about "field" behavior in the lab itself, either through highly contextualized experiments or self-reports. In this case, one might worry about possible spurious correlations caused by demand effects and/or subjects' willingness to avoid cognitive dissonance. Benz and Meier (2008) address part of the above concern by collecting field data about their subjects' behavior in a charitable giving situation prior to conducting a charitable giving experiment in the classroom, but the experiments on which they rely to elicit preferences remain highly contextualized. Barr and Serneels (2009) conduct a Trust game among Ghanaian workers and establish a relationship between individual measures of reciprocity and the observed *aggregate* labor productivity of the firm in which they work. Similarly, Carpenter and Seki (2011) conduct a repeated Public Goods game among Japanese fishermen and show that fishing crews that exhibit higher levels of reciprocity are more productive. Perhaps most similar to the present study, Carpenter and Myers (2010) rely on an

experimental measure of altruism (from a standard Dictator game) and an observational measure of social image concerns within a population of volunteer firefighters and non-volunteer community members to show that both preferences predict the decision to join the volunteer fire service. Finally, Fehr and Leibbrandt (2011) and Leibbrandt (2012) conduct a Public Goods game among Brazilian shrimp catchers and sellers, respectively, and show that more prosocial shrimp catchers are less likely to engage in overextraction, while more prosocial shrimp sellers achieve higher prices for the same goods. While both studies convincingly establish that levels of cooperation in a standard Public Goods experiment can predict field cooperation and economic outcomes, they are not designed, however, to answer the question of which specific preferences account for those general cooperative dispositions.

The present study distinguishes itself from the above literature by eliciting and examining the relative predictive power of all three classes of prosocial motives in a comprehensive fashion. It is also the first to concurrently (i) follow the experimental economics standard of relying on highly decontextualized experiments to elicit individual preferences (ii) link those preferences to *individual* outcomes that were *independently* collected from the field and (iii) examine a real-world public goods like environment in which extrinsic motives play no role in shaping individual behavior.

The rest of the paper proceeds as follows. Section 2 provides some knowledge background on the Wikipedia project and its community of contributors. Section 3 documents the design and implementation of the study. We report the empirical results in section 4. Section 5 provides a discussion of our findings and concludes.

2 Background on Wikipedia

Wikipedia is a free online encyclopedia that is collaboratively edited by volunteers over the Internet. The Wikipedia project originates in Jimmy Wales and Larry Sanger's attempt at creating a traditional, extensively peer-reviewed online encyclopedia called "Nupedia" in March 2000. The goal of Nupedia was to get scholars and experts to volunteer their work and expertise to the project, with the goal of creating a free equivalent of the existing for-profit encyclopedias. Confronted with the difficulty of taking the project off the ground – Nupedia only got 21 articles finalized in its first year – Wales and Sanger eventually released Nupedia's content over the Internet in January 2001 as an open side project, called "Wikipedia", whose original purpose was to feed Nupedia with additional draft articles. Wikipedia quickly overtook Nupedia and became a multiple language popular project of its own, with over 20,000 encyclopedia articles created in its first year and an exponential growth of its content and contributor base since then.

Since 2003, Wikipedia has been operated by the Wikimedia Foundation, a small San Francisco-based non-profit organization, whose role is to pay the bandwidth bills, buy the servers and provide legal defense for the project. The Wikimedia Foundation mostly leverages the capital that it needs to perform this function through donations. It is important to note that while the Foundation is interested in developing technical and social solutions that could support volunteers' editing work, it has never been directly involved in developing Wikipedia's content or managing its community of contributors. This is a matter of principle, and the relationships between the Wikimedia Foundation and the body of engaged Wikipedia contributors have sometimes been notably tense, as some would repeatedly suspect the Foundation of covertly trying to influence the evolution of the project and direct its development.

On the technological side, Wikipedia is based on the wiki system, which allows the reader of any Wikipedia page to modify it easily and rapidly by clicking on an "edit" button. As a result, there exist no limitations *à priori* as to whom can contribute content to the encyclopedia. It is not necessary to create

a Wikipedia account in the website in order to contribute, as this can be done "anonymously". Many regular contributors choose to create a Wikipedia account, however, notably because it gives them access to very useful collaborative editing tools. One prominent example of such a tool is the so-called "watchlist" system, which allows registered users to mark pages of interests and get automatic notices whenever a modification is implemented to them by another contributor. The wiki system archives each and every version of a given page in a revision history, together with the username of the registered contributor who authored the revision. (Contributions made "anonymously" are registered together with the IP address of the computer from which it was performed.) It is customary for contributors to leave a brief summary of their contribution together with the reason why they implemented it upon saving their modifications. This "edit summary" can be read directly from the revision history of any page, which allows contributors to get a very quick sense of each modification and the justification behind its implementation. If necessary (for instance in cases of vandalism), the revision history allows contributors to easily revert a page to one of its previous state.

If they create a Wikipedia account, contributors automatically get a personal user page and a user talk page on the Wikipedia website. Those pages, like virtually every other on Wikipedia, can be edited by anyone. User pages are mostly edited by their owners to post some general information about themselves, their interest in Wikipedia, the articles they helped improve and the like. As collaborations between editors mostly form when they notice that they contribute to the same articles, either through its revision history or the watchlist system (as opposed to randomly scrawling contributors' personal user pages in search for an editor with matching interests), those pages are not crucial to the functioning of Wikipedia. Hence, a significant number of contributors choose to leave them blank. User talk pages, by contrast, are mainly edited by one's fellow editors. They play a very critical role on Wikipedia, as they are used as a convenient place for contributors to communicate with one another, request help, ask questions and coordinate work. Taken together, those technical features explain that even if many individuals do contribute to Wikipedia without having registered an account, the contributions made in this fashion are more likely to be one-offs and, in any case, cannot be much collaborative in terms of content.

The number of contributions made to Wikipedia by registered users follows a strong power law distribution. Skewness of participation characterizes many technology mediated peer production systems. It is not unique to Wikipedia and is also a structural feature of individual contributions to Open Source Software and participation in online message boards. As of 2011, about 200,000 individuals register an account on Wikipedia each month. About 2% of those individuals make 10 contributions or more within their first month, which certainly represents a non negligible influx of new contributors per month in absolute terms. However, only 10% of those early contributors still make one contribution or more within the following year.[9] As a result, the relatively limited body of editors who eventually become engaged and reach the threshold of 100 Wikipedia contributions was still responsible for almost 70% of all the contributions made in 2007 (Kittur et al. 2007). Even within the group of editors who become engaged with the project, individual contribution patterns are still highly heterogeneous. While the vast majority of engaged editors have a few hundred contributions in total, about 5,000 of them made more than 10,000 contributions and about 200 editors have contribution records ranging from 100,000 to 1,000,000 contributions. Overall, the size of the body of active experienced contributors who reached the threshold of 300 contributions is relatively stable since 2007, revolving around 20,000 individuals.

[9] See http://stats.wikimedia.org/EN/TablesWikipediaEN.htm and http://strategy.wikimedia.org/wiki/Editor_Trends_Study/Results, accessed February 2013.

One surprising fact about Wikipedia is the ability of its community of engaged contributors to successfully synthesize into coherent and structured articles their often competing or opposed views about the topics at hand in a civil way. In this respect, it is interesting to note that among subjective topics, the more controversial ones are on average *better* treated in Wikipedia, precisely because they attract attention from a larger and more diverse pool of contributors (Greenstein & Zhu 2012).[10] Reagle (2010) provides a very detailed account of how relationships within the community of engaged editors are generally driven by common behavioral norms that emerged through progressive consensus building as it faced collective action problems. One paradigmatic example of such a norm is the "neutral point of view" policy. It is remarkable that this policy doesn't state that editors should strive to be "neutral" or "objective" while contributing to a given article, but that a "fair" representation of all sides of the dispute should be sought. Conditional on being able to support one's point with reliable secondary sources, this guiding principle has the positive effect of shifting many debates from the question of whether it should be included in the article to the question of *how* it should be included. Another example is the "assume good faith" principle, which exhorts editors to approach others' contributions as being made in good faith and trying to help the project, unless there is specific evidence of malice. When direct discussion fails to resolve disputes among contributors, this is usually achieved by extending the debate to a larger audience, or seeking the mediation of a third party.

Besides the sheer number of contributions that they make to Wikipedia, the body of engaged contributors is thus key to the project, as they often make contributions across topical boundaries in order to curate the content and turn it into a comprehensive resource, help newcomers learn the behavioral norms and attitudes that will allow them to connect with others and make valuable contributions to the project and informally mediate disputes. In this sense, engaged Wikipedia contributors create the public good value of the encyclopedia, and distinguish its contributor base from a broad collection of individuals trying to push their own personal agendas within the site.

One particular class of engaged contributors, the Wikipedia administrators, are in charge of dealing with disruptive editors when good faith discussion and basic explanations about what the goal of the project is fail. To do so, they are entitled with special oversight rights over the encyclopedia that allows them to enforce the behavioral norms of the community, notably by blocking malicious editors, deleting pages that they think have no potential for developing as proper encyclopedic articles and protecting vandalized pages from editing. To obtain those policing rights, those engaged contributors decided to participate in a very competitive peer-review process that would require them to prove through their contribution history that they are valuable contributors who can handle heated debates and make difficult decisions.

3 Design of the study

In this section, we first describe our strategy for measuring social motives among our subjects. We then describe our experimental procedure before reporting on the practical implementation of the experiment.

[10] Controversy on Wikipedia is not limited to "hot" topics such as global warming or the Israeli-Palestinian conflict. One of many mundane examples is the controversy that arose in 2006 around the article on arachnophobia, in which one contributor added the picture of a tarantula. The question of interest was to determine whether the picture had any illustrative value or if it would simply drive potential readers away. A consensus eventually emerged around the idea of replacing the picture by a cartoon illustrating the fear of spiders, but only after several editors had spend hours on the issue, generating around 6,000 words of discussion for an article which is about 1,500 words long.

3. 1 Measuring social preferences

We elicit social motives among our subjects using experimental data from three mostly standard decision problems taken from the literature on social preferences (see, e.g., Fehr & Camerer 2004) coupled with observational evidence. We systematically provide two different measures for each social motive, so that we can check for the consistency of our results. At the beginning of the experiment, subjects are sequentially attributed a role (according to their login order): either participant A or participant B. The assigned role remains the same during the whole experiment. At the end of the experiment, we ask subjects some standard demographic questions about their age, gender, education and salary level, along with an experimentally validated question on risk aversion taken from Dohmen et al. (2011).

(i) Reciprocity motive. Following Fischbacher et al. (2001), we use a modified version the Public Goods game to elicit subjects' reciprocity motive. We start by eliciting subjects' propensity to cooperate in a very standard Public Goods dilemma (see figure 1 which pictures the Public Goods game instructions screen). Subjects play in groups of four with an initial endowment of $10 per player. Each euro invested in the common project by a member of the group yields a return of 0.4 euro to each group member.[11] Subjects have to decide on how much of their $10 they want to invest in the common project. In a second step, we implement the so-called "strategy method" and ask subjects to provide their intended contribution for each possible value (on the scale of integers from 0 to 10) of the average contribution of the three other members. Subjects are told that their actual contribution to the common project will be randomly determined to be either their unconditional contribution from the standard Public Goods game or their conditional contribution decision. *We take the average proportion of the endowment that is conditionally contributed in the conditional Public Goods game as a measure of subjects' reciprocity motive.*

[FIGURES 1 ABOUT HERE]

In order to provide an alternative measure for the reciprocity motive, we also conduct a standard Trust game among our subjects. Each participant A is matched with a participant B, and both players receive a $10 initial endowment. Participant A is the trustor and chooses how much of his endowment is transferred to participant B – the trustee. The trustee receives three times the amount sent by the trustor, and chooses how much is sent back to him. We elicit this decision through the strategy method: for each possible transfer from the trustor (from 1 to 10) the trustee chooses how much will be returned without knowing the trustor's actual choice. *We take the average proportion of the amount received that is returned by the trustee in the Trust game as an alternative measure of subjects' reciprocity motive.*

(ii) Altruistic motive. The Dictator game is certainly experimental economics' workhorse for studying altruistic motives. We thus use a standard Dictator game to elicit this preference among our subjects.[12] Each participant A is matched with a participant B to play as a dictator. The dictator receives a $10

[11] Each subject thus faces the following payoff function: $\pi_i = 10 - \text{contrib}_i + 0{,}4 \sum_{j=1}^{4} \text{contrib}_j$

[12] Note that the measures of altruism that we get from our Dictator games add-up the theoretically distinct "pure altruism" and "warm glow" motives. In this paper, we thus consider the joint effect of those two sub-components of altruism.

endowment, of which he must decide on how much is transferred to participant B. *We take the proportion of the endowment transferred by the dictator as a measure of subjects' altruistic motive.*
As we worry that the standard Dictator game may not capture subjects' altruistic motive if they are incentivized to contribute to Wikipedia out of altruism directed towards their fellow contributors, we provide an alternative measure for this motive by conducting a second Dictator game in which we induce some in-group bias. We do this by telling subjects that they are now matched with another subject who "participates in online collaborative projects such as open source, free software or Wiki-based authoring projects". *We take the proportion of the endowment transferred by the dictator in this directed decision as an alternative measure of subjects' altruistic motive.*

(iii) Social image motive. Social image motives are difficult to measure experimentally – even more so in a decontextualized fashion, that is, out of a given social context. As a result, we rely on the observational data available from Wikipedia in order to elicit this preference. Specifically, we collect the size (in bytes) of the personal user pages of our subjects and use this information to construct a measure of revealed preference for social image within Wikipedia (recall from section 2 that personal user pages do not play an important functional role in Wikipedia and are mainly used to present oneself to the community of contributors). *Separating out regular contributors from Wikipedia administrator, we code as "social signaler" those who have a personal Wikipedia user page whose size (in bytes) is higher than the median in their group, and take this variable as a measure of subjects' social image motive.*
In order to provide an alternative measure for subjects' social image motive, we exploit Wikipedia's main social rewarding practice: the Barnstars system. A Barnstar is a symbolic award constituted of an image accompanied by a personalized message acknowledging some important contribution made to the project by an editor (see figure 3 for an example).[13]

[FIGURE 2 ABOUT HERE]

In theory, anyone can give or receive a Barnstar. This practice, however, remains largely limited to the body of engaged Wikipedia contributors who display relatively impressive contribution records. Barnstars are typically posted on a contributor's talk page. They thus appear within the flow of discussions between this contributor and the rest of the community. After some time, a particular discussion thread is likely to be archived and/or become too long for anyone to easily notice that an award had been given. Some Wikipedia contributors choose to circumvent this by manually moving (some of) their Barnstars to their personal user pages (or some particular subsection of their user page generally labeled their "awards page"), so that they would be durably and prominently displayed for any other editor to see. We take such decisions as revealing a contributor's motive for social image. *From the subsample of subjects who received Barnstars (about 54% of our sample, the vast majority of whom are highly engaged contributors), we code as "social signalers" those who decided to display at least one of those awards on their personal user page, and take this variable as an alternative measure of subjects' social image motive.*

3. 2 Experimental procedures

[13] The Wikipedia "Barnstars" page starts as follows: "It is the custom to reward Wikipedia contributors for hard work and due diligence by awarding them a barnstar. To give the award to someone, just place the image on their talk page (or their awards page), and explain why it was given. If you are sure the barnstar is appropriate, *don't be shy!"* See http://en.wikipedia.org/wiki/Wikipedia:Barnstars, accessed February 2013.

The online implementation of the experiment requires a fully self-contained interface, so that every communication between the subjects and the experimenter has to proceed through the screen. The welcome page of the decision interface provides subjects with general information about the experiment, including the number of sections, expected completion time (about 25 minutes) and how their earnings will be computed. In order to minimize potential demand effects and in-group biases when eliciting subjects' social motives, we were very careful not to present the study as Wikipedia oriented.[14] Importantly, we made it very clear on the introductory screen that subjects would interact with a diverse pool of Internet users.[15]

Subjects are only informed of their earnings in each game at the very end of the experiment. Final payoffs are equal to the earnings from one randomly selected game plus a $10 participation fee (subjects earned on average $20.50 from the experiment). Subjects get paid upon completion of the experiment through an automated PayPal transfer.[16] We only require a valid e-mail address to process the payment. To strengthen the credibility of the payment procedure, we ask subjects to enter the e-mail address that is (or will be) associated with their PayPal account right after the introductory screen of the decision interface. It is important to stress that Wikipedia contributors can be very hostile to monetary rewards. In order to ensure that the experiment is equally incentive compatible for all subjects, we allow them to donate any amount taken from their final earnings to the Wikimedia Foundation and/or the International Committee of the Red Cross – a renowned and general purpose charitable organization, in anticipation of the fact that some subjects might not want to donate to the Wikimedia Foundation – upon completion of the experiment. This possibility was made clear on the welcome screen of the decision interface. It was not possible, however, to commit to donating one's final earnings prior to the study's completion.

All five decisions, followed by the survey, are made successively following a given sequence of screens. The unconditional and conditional Public Goods games are the most cognitively demanding. Accordingly, we always present those two decision problems first to subjects (in this order). As we don't want the Dictator game with induced in-group bias to generate spillover effects on the other decisions, we always maintain both Dictator game decisions in last position. In order to alleviate anchoring effects, we sequentially vary the order in which the standard Dictator game and the directed Dictator game are presented to subject according to their login order. As a result, the standard Trust game was always presented in middle position.

All decisions made by our subjects are anonymous. This is because contrary to the social image motive – which is by definition a *public* social preference – all the preferences that we elicit experimentally are *private* social preferences, meaning that they do not depend on the visibility of one's

[14] The specific language used on the welcome page was as follows: "Our goal is to better understand the dynamics of interactions and behavior in online social spaces. To do so, **we invite internet users with various profiles to fill out an interactive survey**. We very much welcome participation from Wikipedia users!" Our strategy for framing the study as non Wikipedia oriented eventually proved more effective than we had anticipated. When we presented this research project to the Wikimedia Foundation staff, their initial reaction was: "Several people expressed concerns that there was not a clear connection between the contents of the survey and data that would be strategically useful at this time to Wikimedia community members and the Foundation. […] We hope that you will find another suitable outlet to recruit participants for your study. We're happy to answer questions about this decision, and we hope in the future to be able to support other projects you may be working on that are relevant to Wikimedia."

[15] The Wikipedia subjects were matched with a traditional pool of laboratory subjects and with open-source software developers who both previously participated in a similar online experiment.

[16] Such a payment procedure guarantees a fungibility similar to that of cash transfers in lab experiments, as money transferred via PayPal can be readily used for online purchases or easily transferred to one's personal bank account at no cost.

actions to be at work.[17] As we want to elicit social motives in isolation from strategic concerns and learning effects, each game is only played once and we match subjects in each game according to a perfect stranger procedure.

One important methodological concern with the online implementation of the experiment is to guarantee a quick and appropriate understanding of the decision problems when no interaction with the experimenter is possible. We strengthen the internal validity of our online experiment with three distinctive features of the interface. First, we include suggestive flash animations illustrating the written experimental instructions at the bottom of each game's instruction screen (see figure 2 for the example of the standard Public Goods game).[18] Second, the instructions screens are followed by a screen providing some examples of decisions, along with the detailed calculation of the resulting payoffs for each player. These examples are supplemented on the subsequent screen by earnings calculators. On this interactive page, subjects are allowed to test all the hypothetical scenarios they are interested in before making their decisions in the Public Goods and Trust games. In contrast to the illustrative flash animations, the numeric results of each scenario run by a subject in the earnings calculator screens are explicitly displayed. Last, the system provides a quick access to the instructions material at any moment during decision-making. On all screens, including decision-making ones, a "review description" button gives subjects a direct access to the instructions displayed at the beginning of the game. The system also allows participants to navigate at will from one screen to another – until a decision screen has been passed – through the "Previous" and "Next" buttons located at the bottom of each screen (see figure 3 for the example of the conditional Public Goods game decision screen).

[FIGURES 3 ABOUT HERE]

3. 3 Implementation of the experiment

Our main dependent variable of interest is the total number of field contributions that a subject has made to Wikipedia over his history with the project. A Wikipedia contribution, or "edit", is defined as the action of (i) going to a Wikipedia page (ii) hitting the "edit" tab (iii) implementing a modification and (iv) saving the modification. We only recruit from Wikipedia registered users (i.e. individuals who created a Wikipedia account) in order to be able to track subjects' full contribution records.[19]

In order to recruit as representative a sample of the underlying population of Wikipedia contributors as possible, we need to capture the very wide heterogeneity that characterizes registered editors' contributions patterns (see section 2 for some background statistics on this structural feature of

[17] Note that the concept of "social image motive" as it currently stands in the economics literature conflates several motives (e.g. relative social status within a group or relative competence assessment) all of which crucially depend on the visibility of one's actions to be at work. We do not try to distinguish between those in this paper.

[18] The loop of concrete examples displayed in each animation was first randomly determined and then fixed for each game. The same loop is displayed to all subjects without any other numeric information than the subjects' initial endowments. We decided against displaying a purely random sequence of flash animations as it could have introduced uncontrolled and subject specific noise–through, e.g., anchoring on a particular behavior or sequence of events. Our goal with those animations was to illustrate the basic gist of each decision problem in an accessible way while avoiding to prime specific numerical examples and results in subjects' mind.

[19] One might worry about selection effects here. To be sure, this paper does not try to generate results that could be generalizable beyond the population of registered contributors to Wikipedia. In terms of the potential bias induced on our estimates by this selection criterion, we think that insofar as the mere action of registering a Wikipedia account can, on average, be interpreted as a step towards becoming a contributor to the project, then the coefficients on our prosocial motives variables, if they are significant in the true population model, should be biased downwards (as we select on having gone through that step already).

Wikipedia participation, among many other technology mediated peer production systems). To do so, we decide to recruit our subjects from the three following groups:

(i) The cohort of new Wikipedia contributors, defined as all individuals who registered a Wikipedia account within the 30 days prior to the launch of the experiment, irrespective of the number of contributions (if any) that they made. Eligible population = 190,327 subjects.
(ii) The group of engaged Wikipedia contributors, defined as all contributors who made at least 300 contributions to Wikipedia and are still currently active (i.e. they made at least 20 contributions in the last 180 days).[20] Eligible population = 18,989 subjects.
(iii) The group of Wikipedia administrators. Those highly engaged contributors successfully decided to run for a very selective peer-review process, at the end of which they were entitled with special oversight rights over the encyclopedia in order to perform a policing role. They notably can block disruptive users, delete pages that they consider will not develop as proper encyclopedic articles and protect vandalized pages. Eligible population = 1,388 subjects.

We use the Wikipedia banner system as a convenient recruitment device for our experiment. The banner system is prominently used by the Wikimedia Foundation for its annual fundraising and is thus relatively familiar even to non Wikipedia contributors. It is also used by the community of editors for purposes of extended internal communication (e.g., to advertise events and other community initiatives). As a result, the banner system is certainly the most powerful and trusted way of reaching out to a wide and diverse audience within Wikipedia. In coordination with the Wikimedia Foundation staff, we coded this recruitment banner so that it would be displayed at the top of every Wikipedia page for all logged-in eligible users, until he or she decided either to click on it, or to disable it (see Figure 4, which features the recruitment banner).[21]

[FIGURE 4 ABOUT HERE]

Upon clicking on the banner, eligible users were uniquely identified by the system (through their Wikipedia user id number, which then allowed us to collect their entire contribution history to Wikipedia) and redirected to the welcome screen of our experimental economics platform. Within each of the three above-defined experimental groups, our system sequentially allocated subjects to the role of participant A or participant B according to their login order. Those allocated to the role of participant A were in turn sequentially allocated to one of the two possible ordering of the standard and directed Dictator games (in order to alleviate possible anchoring effects). We implemented this procedure both to ensure that we get relatively balanced samples and to randomize the allocation of participants in the role of participant A and participant B. The experiment was launched on December 8th 2011 and the banner recruited 850 subjects in 8 hours (i.e. about 2 complete answers per minute).

4 Results

[20] Note that this definition of an "engaged contributor" corresponds to the community's criteria for being eligible and able to vote in the 2011 elections of the Wikimedia Foundation Board of Trustees. See http://meta.wikimedia.org/wiki/Board_elections/2011/en#Prerequisites_to_candidacy, accessed February 2013.
[21] This was the first (and is still the only) time in the history of the Wikimedia movement that the banner system was left to use by a third-party. Its selective display system remains Wikimedia's most sophisticated one to date.

Our main dependent variable – the number of Wikipedia contributions made by each subject – follows a strong power law distribution. As our dataset is characterized by heteroskedasticity (Breusch-Pagan test: $p<0.001$), we do not present OLS regression tables based on a log-transformation of our dependent variable, as this would induce substantial bias in our estimates (Silva & Tenreyro 2006). As a more cautious approach, we use the negative binomial pseudo-maximum likelihood estimator, which is not affected by this problem.[22] This estimator is appealing because (i) it naturally accounts for the skewness of our data and (ii) the coefficients remain nicely interpreted as semi-elasticities.[23]

We organize the presentation of our results in three steps. We start by presenting some descriptive statistics about our subjects pool, together with a regression analysis of the relationship between demographic characteristics and patterns of contribution to Wikipedia. We then test for the role of altruism, reciprocity and social image as motives for contributing to the Wikipedia project within our sample of regular contributors. In a third step, we focus the analysis on the conceptually distinct class of Wikipedia administrators.

4. 1 Descriptive statistics and analysis

Table 1 provides some descriptive statistics per experimental group on (i) the number of Wikipedia contributions made by our subjects (ii) our measures of social motives and (iii) our demographic variables. Overall, we recruited 149 subjects from the cohort of new Wikipedia contributors, 566 from the group of engaged Wikipedia contributors and 120 from the group of Wikipedia administrators. Because the data used to calculate the eligibility metrics was missing for some users in the Wikipedia API, 15 Wikipedia contributors were displayed the recruitment banner and participated in the experiment while not being formally eligible to do so. As we are equally able to track the contribution records of those subjects, we also include them in our analysis.[24]

[TABLE 1 ABOUT HERE]

For each experimental group, figure 5 compares the distribution of the number of Wikipedia contributions for the whole sample of eligible contributors against our sub-sample of participants. Focusing on the groups of engaged Wikipedia contributors and Wikipedia administrators, we can see that the distribution of the number of Wikipedia contributions for our sub-samples of subjects closely mirror those of the reference groups. We do seem to capture contributors with higher contribution records on average, however, as we can see from both distributions being slightly skewed to the right. We reach a similar conclusion when we focus on the cohort of new Wikipedia contributors. Out of 149 subjects, 62% have never made any contribution to Wikipedia (as opposed to 73% in the reference group) and 27% made between 1 and 10 contributions (as opposed to 25% in the reference group). 11%

[22] See Wooldridge (2002, chapter 19). Log-linearizing our dependent variable to run OLS regressions yields estimates that are higher in magnitude than the ones presented in this paper. The empirical conclusions remain qualitatively similar, however (tables available from the authors upon request).

[23] An alternative estimator that has similar properties is the Poisson pseudo-maximum likelihood estimator. One limitation of this estimator, however, is that it does not allow for overdispersion (which is a feature of our data, likelihood ratio test: $p<0.001$). The negative binomial estimator is more flexible and estimates the form of the dispersion as an additional parameter.

[24] All of the results presented in this paper remain unaffected if we leave those 15 subjects out of the analysis.

of our new contributors, however, are already highly engaged with Wikipedia and made more than 10 – and up to 273 – contributions to Wikipedia (as opposed to 2% in the reference group).[25]

[FIGURE 5 ABOUT HERE]

Another way to look at how representative of the overall population of Wikipedia registered editors our sample of subjects might be is to pool them all together and compare their demographic characteristics against those of the 5,073 registered editors who took part in the 2011 Wikimedia editor survey. Designed by the Wikimedia Foundation, this survey was precisely implemented so as to get as representative a picture as possible of the profiles of Wikipedia editors.[26] Similar to the present study, it was advertized through a Wikipedia banner. It ran for 7 days over the whole population of registered Wikipedia editors. Table 2 compares the commonly available demographic information in both studies. It appears that demographic characteristics between both samples are very similar. Contrary to the popular perception that most Wikipedia contributors are high school students, we find that they are on average much older (33 years old with 48% of the population being above 29 in our study versus 32 years old with 47% being above 29 in the Wikimedia editor survey) and more educated (63% have finished college and 28% have a Master's or a PhD degree in our study versus 61 and 26% in the Wikimedia editor survey, respectively). Consistent with the common perception, however, we find the population of contributors to be predominantly male (90% in our study versus 89% in the Wikimedia editor survey). Taken together, we interpret the above evidence as suggesting that our sample of Wikipedia subjects is representative of the diversity of contribution patterns and demographic profiles found on Wikipedia.

[TABLE 2 ABOUT HERE]

We end this section by presenting a regression analysis of the effect of subjects' demographic characteristics on the number of contributions that they make to Wikipedia (see table 3). Column (1) focuses on the group of regular contributors (as opposed to Wikipedia administrators). The model globally confirms our above qualitative observations: being one year older is on average associated with a 1.7% increase in the number of Wikipedia contributions, while moving from a high school education to getting a Master's degree is associated with a 26% increase. Being a female, however, is associated with a 44% decrease in the number of contributions made to Wikipedia. The coefficient on the salary level variable is very close to zero and not statistically significant. This result is surprising, as it suggests that subjects' opportunity cost of time does not have any significant impact on their willingness to contribute to Wikipedia. Finally, the effect of risk aversion seems somewhat counterintuitive: moving from generally being "unwilling to take risks" to being "fully prepared to take risks" is actually associated with a 43% *decrease* in the number of Wikipedia contributions.

Those average effects conceal an interesting underlying heterogeneity within our population of subjects, however. In columns (2) and (3) we divide our sample of regular contributors in two equal parts according to the median of the number of contributions that they made to Wikipedia (i.e. 1905 contributions, which already represents a rather impressive contribution record) and run the exact same regressions as in model (1) for both sub-populations. We can see that while the effect of our demographic variables remains qualitatively the same within the group of new to engaged Wikipedia

[25] A Kolmogorov-Smirnov test of the equality of distribution functions confirms this conclusion at $p<0.001$ in all three experimental groups.

[26] See http://meta.wikimedia.org/wiki/Editor_Survey_2011, accessed February 2013.

editors, none of those variables reliably predict the contribution patterns of those subjects who are already "super contributors" to the project. The trajectories of the highest contributors to Wikipedia thus seem very difficult to foresee, even using those very standard covariates. In all of our subsequent tests, we will therefore systematically check for such heterogeneous effects between both sub-populations of new to engaged Wikipedia contributors on the one hand, and "super contributors" on the other (thereafter denoted as the "below median" and "above median" groups for simplicity).

Last, column (5) of table 3 presents the impact of our demographic variables on Wikipedia participation within the group of Wikipedia administrators. In this group, being one year older is on average associated with a 1.6% increase in the number of Wikipedia contributions made. The coefficient on the salary level variable achieves statistical significance in this particular group: out of 9 possible revenue categories, moving from one to the next is associated with a 5.2% decrease in the number of Wikipedia contributions made. This is interesting, as it suggests that the opportunity cost of time has a negative impact on Wikipedia participation only within the group that typically features the most extreme contribution patterns to Wikipedia.

[TABLE 3 ABOUT HERE]

With those results in mind, we now turn to our theoretical question of interest, and investigate which prosocial motives, altruism, reciprocity or social image, better account for subject's willingness to contribute to Wikipedia in the field. The next section focuses on the prosocial foundations of cooperation among regular contributors, while the following one focuses the analysis on the group of Wikipedia administrators, which we consider as a conceptually distinct class of contributors.

4. 2 Prosocial foundations of cooperation: regular Wikipedia contributors

We test for the role of each class of social motive on subjects' willingness to sustain their contributions to Wikipedia by including our measures of altruism, reciprocity and social image motives in turn in regressions that include our above demographic variables as covariates. Having two alternative measures for each class of social motive allows us to check for the consistency of the results that we get. Table 4 tests for the predictive power of both experimental measures of general and directed altruism in the Dictator game. We can see that no statistically significant relationship appears with field contributions. This is true irrespective of whether we consider the whole sample of regular subjects (columns (1) and (2)) or, as in table 3, check for potential heterogeneous effects by separating the sample of contributors in two sub-populations according to the median of their number of Wikipedia contributions (columns (3) to (6)).

[TABLE 4 ABOUT HERE]

The picture is completely different when we turn to table 5, which tests for the role of the reciprocity motive. According to our estimates, moving from no reciprocity to full reciprocity in the conditional Public Goods game and in the standard Trust game is associated with a significant 46% and 56% increase in the number of Wikipedia contributions, respectively. Similar to what we find in table 3 in terms of the effect of demographic characteristics, those average estimates conceal an interesting

heterogeneity within the population of contributors, however. Columns (3) and (4) reveal that the association between both experimental measures of reciprocity and Wikipedia contributions is much higher in magnitude and highly statistically significant in the below median group, while it is insignificant in the above median group (columns (5) and (6)). Focusing on the below median group, moving from no reciprocity to full reciprocity is associated with a 122% and 211% increase in the number of Wikipedia contributions, depending on the experimental measure of reciprocity that we consider. Those results indicate that subjects' willingness to sustain their contributions to Wikipedia is related to their taste for reciprocity as opposed to altruism. Interestingly, however, while reciprocity appears as the major *private* social preference associated with the trajectory of Wikipedia users from a non-contributor to a regular contributor, this preference does not seem to continue to predict the trajectories of the highest contributors to Wikipedia. We now rely on observational data to investigate the role of social image motives in subjects' willingness to contribute to the Wikipedia project.

[TABLE 5 ABOUT HERE]

Within our sample of regular subjects, we code as "social signalers" those who have a personal Wikipedia user page whose size (in bytes) is higher than the median in the sample. Alternatively, from the sub-sample of subjects who received social awards – or Barnstars – from other Wikipedia contributors (i.e. 456 subjects, representing 54% of our total sample), we code as "social signalers" those who decided to advertize at least one of those awards on their personal user page.[27] According to this measure, 54% of Barnstars receivers reveal a preference for social image. Importantly (and almost by definition), 81% of Barnstars receivers in the sample of regular subjects have contribution records that place them in the above median group. Therefore, one limitation of this variable is that it will mainly tell us about the role of social image motives within the group of highest contributors to Wikipedia. As we expect subjects who receive more Barnstars to have a higher probability of exhibiting one of them on their personal user page (at least in a statistical sense), and as the total number of Barnstars received should be highly correlated with the number of Wikipedia contributions made, we include the total number of Barnstars received as a control in all the regressions that rely on this measure of social image to avoid potential spurious correlations.

Table 6 presents the results of those estimations for all regular contributors. We see in column (1) that subjects who reveal a preference for social image by having a relatively larger Wikipedia user page make on average 269% more contributions to Wikipedia. This highly statistically significant result confirms the hypothesis that those who care relatively more about their social image within the community of editors also contribute more to the Wikipedia project.

To check for heterogeneous effects, we also run the exact same regression as in model (1) separately in the below and above median groups. In agreement with what we found in the case of reciprocity, we can see from column (2) that the coefficient on social image increases by 38% in the below median group and remains highly statistically significant. However, contrary to what we found in the case of reciprocity, a revealed preference for social image continues to be significantly associated with the number of contributions made to Wikipedia even within the group of highest contributors. Indeed, in the above median group, social signalers make on average 30% more contributions to the Wikipedia project (column (3)). The magnitude and significance of this estimate is confirmed when we rely on our Barnstars data to construct an alternative indicator of taste for social image for highly engaged

[27] We also tried using the *proportion* of received Barnstars that subjects decided to manually move to their personal user pages as an alternative indicator of their social image motive. The results were unaffected (table available from the authors upon request).

contributors, as we obtain that social signalers make on average 33% more contributions to Wikipedia by this measure (see column (4)).

[TABLE 6 ABOUT HERE]

So far, we have established that our measures of reciprocity and social image – but not altruism – are reliably associated with the trajectory of Wikipedia users from a non-contributor to a regular contributor. Unlike reciprocity, however, a taste for social image continues to correlate with the number of field contributions made by our subjects even within the group of highest contributors to Wikipedia.

Building upon this result, a natural question to ask is that of the nature of the interaction (if any) between reciprocity and social image as other-regarding motives for contributing to Wikipedia. We answer this question by investigating whether our experimental measures of reciprocity predict the number of contributions that our subjects make to Wikipedia differentially, depending on whether they reveal a concern for their social image within the Wikipedia community or not. To achieve this goal, we re-estimate our coefficients on the reciprocity motive separately for social signalers and non social signalers.

Focusing on column (1) and (2), we can see that, irrespective of the experimental measure of reciprocity that we consider, the predictive power of this preference on the number of field contributions to Wikipedia appears as concentrated within the group of non social signalers, that is contributors who do not reveal a relatively high preference for social image within Wikipedia. Restricting the sample to the below median group and running the same regressions reinforces this conclusion. We can see from columns (3) and (4) that the coefficients on reciprocity in the group of social signalers are positive, but remain statistically insignificant. The coefficients on reciprocity in the group of non social signalers, by contrast, remain highly statistically significant and increase by 66% and 10%, respectively.

Turning our attention to columns (5), (6), (7) and (8), we obtain a picture that is consistent with that of table 5, columns (5) and (6), in which we could not find anymore any statistically significant association between reciprocity and the number of contributions made to Wikipedia within the group of super contributors. In fact, differentiating between social signalers and non social signalers within this group reveals a surprising *negative* correlation between reciprocity preferences and the extent of Wikipedia participation for those highly engaged contributors who are social signalers (with the coefficient being marginally significant in 2 out of 4 cases). By contrast, the coefficients on reciprocity for the non social signalers, if they are not statistically significant, are all positive.

At the end of the day, what those estimations suggest is that both the reciprocity and social image motives are strongly associated with individuals' willingness to sustain cooperation in a real-world public goods like environment such as Wikipedia, but that they seem to be substitutable rather than complementary motivational drivers (i.e. both motives are at play, but in different subsets of the population of contributors).

[TABLE 7 ABOUT HERE]

4. 3 Prosocial foundation of cooperation: Wikipedia administrators

This section replicates the above analysis and discusses the link between prosocial preferences and patterns of contributions to Wikipedia within the group of Wikipedia administrators, a distinct, high

status class of highly engaged contributors who successfully opted-in a very competitive peer review process at the end of which they were granted with special oversight rights over the encyclopedia. Those contributors are in charge of enforcing the behavioral rule and standards of the Wikipedia community when basic communication between contributors fails at achieving a cooperative outcome. To perform their policing and curating role within the community, Wikipedia administrators notably have the ability to block disruptive users, delete the Wikipedia pages that do not have the potential to develop as proper encyclopedic articles and prevent vandalized pages from being edited by certain groups of contributors.

The last two columns of table 4 above present our results on altruism for the group of Wikipedia administrators. Consistent with what we find for the group of regular contributors, we see no statistically significant relationship between either measure of general and directed altruism from the Dictator game and Wikipedia administrators' patterns of contribution.

The picture is different when we turn to the last two columns of table 5, however, which investigate the predictive power on Wikipedia participation of our experimental measures of reciprocity from the conditional Public Goods and Trust games. Within the group of Wikipedia administrators, we find consistent evidence that reciprocity motives are actually negatively associated with the extent of Wikipedia participation. This pattern is statistically significant at the 5% level for both experimental measures. Within the group of Wikipedia administrators, moving from full reciprocity to no reciprocity is associated with a 88% and 169% *decrease* in the number of Wikipedia contributions, depending on the experimental measure of reciprocity that we consider.

We now turn to our observational data to investigate the association between social image motives and Wikipedia administrators' patterns of contribution to the project. One preliminary observation is that requesting Wikipedia adminship implies going through a long and costly peer review process. Referring to his contribution history, the candidate has to convince the community of contributors that he is capable of undertaking this responsibility and typically has to achieve a very high proportion of positive comments on his request to succeed. Hence, the very fact of being willing to go through the process required to get the high status position that goes along with administrator rights can be interpreted in itself as an indicator of taste for social image. Still, we can see from the last two columns of table 6 that our social image measures continue to be positively associated with a higher number of field contributions made to Wikipedia within this group. The evidence is less strong than within the group of regular contributors, however. Dividing our sample of Wikipedia administrators in two groups according to the median size of their personal Wikipedia user page yields a statistically significant result: those who reveal a relatively higher taste for social image make on average 42% more contributions to the Wikipedia project (column (4)). The coefficient does not achieve statistical significance when we compute our indicator of taste for social image using the Barnstars data, however (column (5)).

With those results in mind, we investigate the nature of the interaction between reciprocity and social image in the group of Wikipedia administrators following the same strategy as for regular contributors (see table 8). We observe that reciprocity preferences are strongly negatively associated with Wikipedia participation, irrespective of whether we consider the sub-group of administrators who reveal a relatively lower or relatively higher taste for social image. The individual coefficients reach strong statistical significance in 5 out of 8 cases, despite our reduced sample size for this group. This result tends to reinforce the conclusion from table 5, in which we uncovered a strong negative relationship between administrators' reciprocity preference and their patterns of contribution to Wikipedia.

[TABLE 8 ABOUT HERE]

We end this section by digging further into the finding that Wikipedia administrators who contribute relatively more to Wikipedia are actually less prosocial on average. This relationship may be related to the fact that those engaged contributors self-selected into performing a policing role within the community of editors. Indeed, the task of any system operator is to keep his system secure, up and running, which often involves dealing with a very large number of potentially malicious users and fixing numerous "bugs". Hence, we hypothesize that two (non exclusive) mechanisms could be at work. First, those engaged contributors may self-select into being administrators because one *needs* to be assuming that people are inclined to hurt the system and see them as potential threats in the first place if he wants to be efficient in his task (i.e. a self-selection mechanism based on having a "thick skin"). In this respect, system administrators should feel responsibility towards the system, not the people, which could explain the negative correlation between their activity and prosociality levels. Second, by being exposed to many malicious users, Wikipedia administrators may end up developing low levels of trust towards anonymous strangers (i.e. develop a "system operator syndrome" based on a learning mechanism).

While we have no way to tease those mechanisms out, we can test them jointly in a direct way by relying on the data on trusting behavior that we collected as a byproduct of our measure of reciprocity based on the Trust game. To do so, *we take the proportion of the endowment that trustors decided to send to trustees in the Trust game as an experimental measure of subjects' level of trust towards anonymous strangers.* We then test for the predictive power of this experimental measure of general trust on the patterns of contributions of our Wikipedia administrators (see table 9). As a first step to the analysis, we first verify that trust is not associated with the contribution patterns of regular contributors, as we have no reason to expect that this should be the case in theory. We can verify from columns (1) to (3) that the coefficients on trust for regular contributor are statistically insignificant and close to zero. This is true irrespective of whether we consider the whole sample of regular subjects or check for potential heterogeneous effects by looking at the below median and above median groups separately.

As hypothesized, the picture is different when we focus on the sample of Wikipedia administrators. In this group, moving from full trust to no trust in strangers is significantly associated with a 107% increase in Wikipedia activity (see column (4)). In order to cross validate this result, we specifically collect the paradigmatic administrative activity types of our subjects – number of users blocked, number of pages deleted and number of pages protected from editing – and test for the predictive power of our experimental measure of trust in strangers on the extent to which they engage in those policing activities. As we can we from columns (5) to (7), moving from full trust to no trust in the experiment is associated with a 173% reduction in the total number of users blocked from editing, a 107% reduction in the number of pages deleted, and a 87% reduction in the number of pages protected from editing, with the effect being statistically significant in 2 out of 3 cases.

As a final piece of evidence, we returned to our subjects 6 months after the completion of the experiment (i.e. in July 2012) and asked them to tell us about the fraction of their working time on Wikipedia that they typically spent on activities that administrators only can perform (e.g. deleting and protecting pages, blocking and unblocking users etc.), as opposed to regular contribution activities. We received an answer from 58 Wikipedia administrators out of 120 in the original sample. Column (8) of table 9 presents an OLS estimate of the relationship between trust in anonymous strangers and the fraction of their working time on Wikipedia that those administrators reported dedicating to administrative activities. Despite the small sample size, moving from full trust to no trust in the Trust

game is significantly associated with a 3.7 points decrease in the proportion of time dedicated to admin activities. Out of a 10 points scale, this estimate corresponds to a 1.34 standard deviation decrease.

[TABLE 9 ABOUT HERE]

5 Discussion

The results of our field study can be summarized as follows:

For regular contributors:

i. Reciprocity and social image – but not altruism – clearly appear as underlying social motives that are associated with the trajectory of Wikipedia users from a non-contributor to an engaged contributor.
ii. Reciprocity and social image seem to be substitutes rather than complementary motivational drivers (i.e. each motive is at play, but in different subsets of the population of contributors).
iii. A taste for reciprocity does not continue to be associated with the trajectory of those Wikipedia users who become super-contributor, while a taste for social image does.

For Wikipedia administrators:

i. There is some evidence that a higher taste for social image continues to be associated with higher contributions levels to Wikipedia, even within the high status group of Wikipedia administrators.
ii. Reciprocity preferences are significantly negatively associated with the extent of participation within this group. We posit that this relationship may be explained by the fact that those engaged contributors self-selected into "holding the stick of the community" which implies that they either need to see non established users as potential threats or that they may develop low levels of general trust as a result of being exposed to many malicious users.
iii. We test for the above mechanism by exploiting our experimental measure of trust towards anonymous strangers. We find that less trusting administrators are significantly more active, more likely to block other users from editing, more likely to delete Wikipedia pages and dedicate a higher proportion of their working time on Wikipedia to admin related activities.

Our results are in striking agreement with the findings of the single other related study of Wikipedia in the economics literature. Focusing on the Chinese Wikipedia, Zhang and Zhu (2011) find that after an exogenous reduction in the group size of contributors (i.e. a block of Wikipedia that only affected mainland China), regular contributors who were not affected by the block decreased their contributions by 42.8% on average. The authors hypothesize that their findings might be due to what they call "social effects", that is social benefits that would accrue to contributors as the size of their group grows. Our results support and precise their hypothesis, as models based on reciprocity and social image motives would both predict an increase in individual contributions following an increase in the size of the group of contributors, while models based on altruism would either predict no impact (in the case of warm-glow) or even a decrease (in the case of pure altruism) in individual contributions.

Beyond the economics literature, our results are also in line with the computer science literature on Wikipedia. For instance, Choi et al. (2010) conduct a quantitative study of what they call "socialization tactics on Wikipedia", whereby engaged contributors use various strategies to reach out to newly

registered contributors and get them to contribute. They find that sending welcome messages, providing assistance and making constructive criticisms significantly increase the likelihood of engagement. Similarly, Halfaker et al. (2011) find that new Wikipedia contributors who see their early contributions reverted by more senior contributors are very likely to stop contributing, but that providing some feedback on the reason for the revert significantly counteracts this effect. One conclusion of their study is that "the more curmudgeonly old-timers should be kept away from newcomers until they have gained some experience in the system."

Our results concerning the group of Wikipedia administrators could be interpreted in light of the above observation. The less trusting administrators, who typically spend a relatively higher fraction of their working time on Wikipedia performing administrative actions and are found to be more likely to block users and delete pages, could be considered as archetypal examples of such surly old-timers. Such an interpretation would echo a current debate in the Wikipedia community about whether "newbie biting" could be one of the factors behind Wikipedia's growing difficulties at attracting more contributors willing to sustain high contribution levels.[28] It is important to note, however, that our data does not allow us to confirm this interpretation. As mentioned earlier, it could also be that an efficient Wikipedia administrator needs be relatively suspicious of anonymous strangers in order to successfully protect the project from non established users who could be inclined to hurt it.

At the end of the day, our results have potentially important implications for practitioners who seek to leverage intrinsic motivations to promote Internet-mediated voluntary cooperation for the provision of global public goods. If anything, from Wikipedia to Open Source Software, the impressive success of peer production in the last 20 years is an indication that intrinsic motivations generally construed (including, but probably not limited to, prosocial motivations) can be very powerful at incentivizing work. Wikipedia is a textbook case for the peer production model, as well as a striking success story. How much this model will continue to scale-up probably depends on how good practitioners will be at efficiently designing large scale human interaction systems that motivate voluntary participation. Our findings suggest that to maximize individual contributions, some special emphasis should be put on the human interactivity side of those systems coupled with some public recognition mechanisms, which will notably continue to incentivize the highest potential contributors.

6 Conclusion

Peer production is certainly the most significant organizational innovation that has emerged from Internet-mediated social practices (Benkler 2013). The distinctive features of this emerging production model – voluntary self-assignment of work and successful large scale coordination in the absence of price signals, pre-specified design rule or formal leadership – make it difficult to understand fully through the assumptions of standard economic theory.

Taking Wikipedia as one paradigmatic example of peer production in which monetary incentives play a negligible role in shaping individual behavior, we build upon the theory of voluntary cooperation in public goods like environments to provide an account of the prosocial motives that could explain individuals' willingness to contribute time and effort towards the provision of non excludable goods. By doing so, we also provide the first comprehensive field test of the existing economic theories of prosocial motives for contributing to real-world public goods. We elicit the social preferences of a representative sample of 850 Wikipedia contributors with an online experiment coupled with observational data and

[28] See http://strategy.wikimedia.org/wiki/Editor_Trends_Study/Results

test for their predictive power over records of contributions to the Wikipedia project. We find sizeable relationships between individuals' prosocial motivations and their patterns of contributions in this peer production economy.

The results of this field study have important theoretical implications, as they strongly support the models of voluntary provision of public goods based on reciprocity and social image motives, but not those based on altruistic motives. In this respect, it is reinsuring to note that this overarching conclusion is strongly consistent with the results of the extensive literature from the lab that has tried to test for the role of those three classes of social motives in people's willingness to sustain cooperation in repeated public goods experiments (see footnote 8).

Of course, while economic theory typically assumes that individual preferences are fixed, our experiment does not allow us to tell whether the preferences that we elicit actually cause the subsequent patterns of contributions that we observe or whether they merely evolved as a result of Wikipedia participation. It is important to note in this respect that when speak of the "predictive power" of our experimental measures of social motives on field behavior, we mean it in a very precise sense that has to do with the external or "ecological" validity of experimental measures that have been used extensively in the lab to test economic theory. Consistent with the fact that most highly engaged Wikipedia contributors started to contribute intensely from the very start of their career (Panciera et al. 2009), we believe, however, that some individual preferences do have a causal impact on subsequent contribution patterns.

We are, of course, only beginning to uncover the nature of the intrinsic motives that drive individuals to voluntarily sustain cooperation in the field. These motives are likely to be diverse. Much more field work needs to be done to see if the literature will be able to identify some general underlying preferences that would be systematically associated with sustained patterns of contribution to real-world public goods, irrespective of the context in which such contributions take place. It could also be, however, that the motives that drive contributions highly depend on the nature of the public good considered, which could in turn explain some of the contradicting laboratory results in the literature (see Vesterlund (2012) and Ostrom (1990)). Although the Internet is a rather specific field of study, we suggest that there is increasing scope for learning from an online approach coupling the tools of experimental economics with computational social science techniques. This is true in the sense that the Internet allows to run experiments eliciting individual preference parameters from large and diverse populations, and to connect those preferences to very detailed observational data on individual field behavior.

References

Anderson, S.P., Goeree, J.K. & Holt, C.A., 1998. A theoretical analysis of altruism and decision error in public goods games. *Journal of Public Economics*, 70(2), pp.297–323.

Andreoni, J., 1995. Cooperation in Public-Goods Experiments: Kindness or Confusion? *The American Economic Review*, 85(4), pp.891–904.

Andreoni, J., 1989. Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. *Journal of Political Economy*, 97(6), pp.1447–1458.

Andreoni, J., 1990. Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal*, 100(401), pp.464–477.

Andreoni, J. & Bernheim, B.D., 2009. Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects. *Econometrica*, 77(5), pp.1607–1636.

Andreoni, J. & Miller, J., 2002. Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism. *Econometrica*, 70(2), pp.737–753.

Andreoni, J. & Petrie, R., 2004. Public goods experiments without confidentiality: a glimpse into fund-raising. *Journal of Public Economics*, 88(7–8), pp.1605–1623.

Andreoni, J., Rao, J.M. & Trachtman, H., 2011. Avoiding The Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving. *Working Paper*.

Ariely, D., Bracha, A. & Meier, S., 2009. Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially. *American Economic Review*, 99(1), pp.544–55.

Barr, A. & Serneels, P., 2009. Reciprocity in the workplace. *Experimental Economics*, 12(1), pp.99–112.

Bénabou, R. & Tirole, J., 2006. Incentives and Prosocial Behavior. *The American Economic Review*, 96(5), pp.1652–1678.

Benkler, Y., 2002. Coase's Penguin, or, Linux and "The Nature of the Firm." *The Yale Law Journal*, 112(3), pp.369–446.

Benkler, Y., 2013. Peer Production and Cooperation. In *Handbook on the Economics of the Internet, J. M. Bauer & M. Latzer (eds.)*. Cheltenham and Northampton, Edward Elgar.

Benkler, Y., 2006. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*, Yale University Press.

Benz, M. & Meier, S., 2008. Do people behave in experiments as in the field?—evidence from donations. *Experimental Economics*, 11(3), pp.268–281.

Burlando, R. & Guala, F., 2005. Heterogeneous Agents in Public Goods Experiments. *Experimental Economics*, 8(1), pp.35–54.

Carpenter, J. & Myers, C.K., 2010. Why volunteer? Evidence on the role of altruism, image, and incentives. *Journal of Public Economics*, 94(11–12), pp.911–920.

Carpenter, J. & Seki, E., 2011. Do social preferences increase productivity? Field experimental evidence from fishermen in Toyama Bay. *Economic Inquiry*, 49(2), pp.612–630.

Charness, G. & Yang, C.-L., 2007. Endogenous Group Formation and Public Goods Provision: Exclusion, Exit, Mergers, and Redemption. *Working Paper*.

Choi, B. et al., 2010. Socialization Tactics in Wikipedia and Their Effects. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. CSCW '10. New York, NY, USA: ACM, pp. 107–116.

Cinyabuguma, M., Page, T. & Putterman, L., 2005. Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics*, 89(8), pp.1421–1435.

DellaVigna, S., List, J.A. & Malmendier, U., 2012. Testing for Altruism and Social Pressure in Charitable Giving. *The Quarterly Journal of Economics*, 127(1), pp.1–56.

Dohmen, T. et al., 2011. Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences. *Journal of the European Economic Association*, 9(3), pp.522–550.

Dufwenberg, M. & Kirchsteiger, G., 2004. A theory of sequential reciprocity. *Games and Economic Behavior*, 47(2), pp.268–298.

Ellingsen, T. & Johannesson, M., 2011. Conspicuous generosity. *Journal of Public Economics*, 95(9–10), pp.1131–1143.

Ellingsen, T. & Johannesson, M., 2008. Pride and Prejudice: The Human Side of Incentive Theory. *The American Economic Review*, 98(3), pp.990–1008.

Falk, A. & Fischbacher, U., 2006. A theory of reciprocity. *Games and Economic Behavior*, 54(2), pp.293–315.

Fehr, E. & Camerer, C.F., 2004. Measuring Social Norms and Preferences Using Experimental Games: A Guide for Social Scientists. *Foundations of Human Sociality*, 1(9), pp.55–96.

Fehr, E. & Leibbrandt, A., 2011. A field study on cooperativeness and impatience in the Tragedy of the Commons. *Journal of Public Economics*, 95(9–10), pp.1144–1155.

Fischbacher, U. & Gächter, S., 2010. Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments. *The American Economic Review*, 100(1), pp.541–556.

Fischbacher, U., Gächter, S. & Fehr, E., 2001. Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3), pp.397–404.

Gächter, S. & Thöni, C., 2005. Social learning and voluntary cooperation among like-minded people. *Journal of the European Economic Association*, 3(2-3), pp.303–314.

Giles, J., 2005. Internet encyclopaedias go head to head. *Nature*, 438(7070), pp.900–901.

Goeree, J.K., Holt, C.A. & Laury, S.K., 2002. Private costs and public benefits: unraveling the effects of altruism and noisy behavior. *Journal of Public Economics*, 83(2), pp.255–276.

Greenstein, S. & Zhu, F., 2012. Collective Intelligence and Neutral Point of View: The Case of Wikipedia. *National Bureau of Economic Research Working Paper*.

Gunnthorsdottir, A., Houser, D. & McCabe, K., 2007. Disposition, history and contributions in public goods experiments. *Journal of Economic Behavior & Organization*, 62(2), pp.304–315.

Halfaker, A., Kittur, A. & Riedl, J., 2011. Don'T Bite the Newbies: How Reverts Affect the Quantity and Quality of Wikipedia Work. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*. WikiSym '11. New York, NY, USA: ACM, pp. 163–172.

Hergueux, J. & Jacquemet, N., 2013. Social Preferences in the Online Laboratory: A Randomized Experiment. *Working Paper*.

Holländer, H., 1990. A Social Exchange Approach to Voluntary Cooperation. *The American Economic Review*, 80(5), pp.1157–1167.

Karlan, D.S., 2005. Using Experimental Economics to Measure Social Capital and Predict Financial Decisions. *The American Economic Review*, 95(5), pp.1688–1699.

Kittur, A. et al., 2007. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *Proceeding of the CIGCHI Conference on Human Factors in Computing Systems*.

Laury, S.K. & Taylor, L.O., 2008. Altruism spillovers: Are behaviors in context-free experiments predictive of altruism toward a naturally occurring public good? *Journal of Economic Behavior & Organization*, 65(1), pp.9–29.

Lazer, D. et al., 2009. Life in the network: the coming age of computational social science. *Science (New York, N.Y.)*, 323(5915), pp.721–723.

Leibbrandt, A., 2012. Are social preferences related to market performance? *Experimental Economics*, 15(4), pp.589–603.

Lerner, J. & Tirole, J., 2002. Some Simple Economics of Open Source. *The Journal of Industrial Economics*, 50(2), pp.197–234.

Oliveira, A.D., Croson, R.T.A. & Eckel, C.C., 2009a. Are Preferences Stable Across Domains? An Experimental Investigation of Social Preferences in the Field. *Working Paper*.

Oliveira, A.D., Croson, R.T.A. & Eckel, C.C., 2009b. One Bad Apple: Uncertainty and Heterogeneity in Public Good Provision. *Working Paper*.

Ostrom, E., 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press.

Page, T., Putterman, L. & Unel, B., 2005. Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency. *The Economic Journal*, 115(506), pp.1032–1053.

Palfrey, T.R. & Prisbrey, J.E., 1997. Anomalous Behavior in Public Goods Experiments: How Much and Why? *The American Economic Review*, 87(5), pp.829–846.

Panciera, K., Halfaker, A. & Terveen, L., 2009. Wikipedians Are Born, Not Made: A Study of Power Editors on Wikipedia. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work*. GROUP '09. New York, NY, USA: ACM, pp. 51–60.

Rabin, M., 1993. Incorporating Fairness into Game Theory and Economics. *The American Economic Review*, 83(5), pp.1281–1302.

Reagle, Joseph M., 2010. *Good Faith Collaboration: The Culture of Wikipedia*, MIT Press.

Rege, M. & Telle, K., 2004. The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics*, 88(7–8), pp.1625–1644.

Silva, J.M.C.S. & Tenreyro, S., 2006. The Log of Gravity. *Review of Economics and Statistics*, 88(4), pp.641–658.

Vesterlund, Lise, 2012. Voluntary Giving to Public Goods: Moving Beyond the Linear VCM. *Working Paper*.

Vesterlund, Lise, Wilhelm, Mark & Xie, Huan, 2009. Motives for charitable giving. *Working Paper*.

Wooldridge, J.M., 2002. *Econometric Analysis of Cross Section and Panel Data*, MIT Press.

Zhang, X. & Zhu, F., 2011. Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia. *The American Economic Review*, 101(4), pp.1601–1615.

Tables and figures

Figure 1. The instruction screen of the Public Goods game

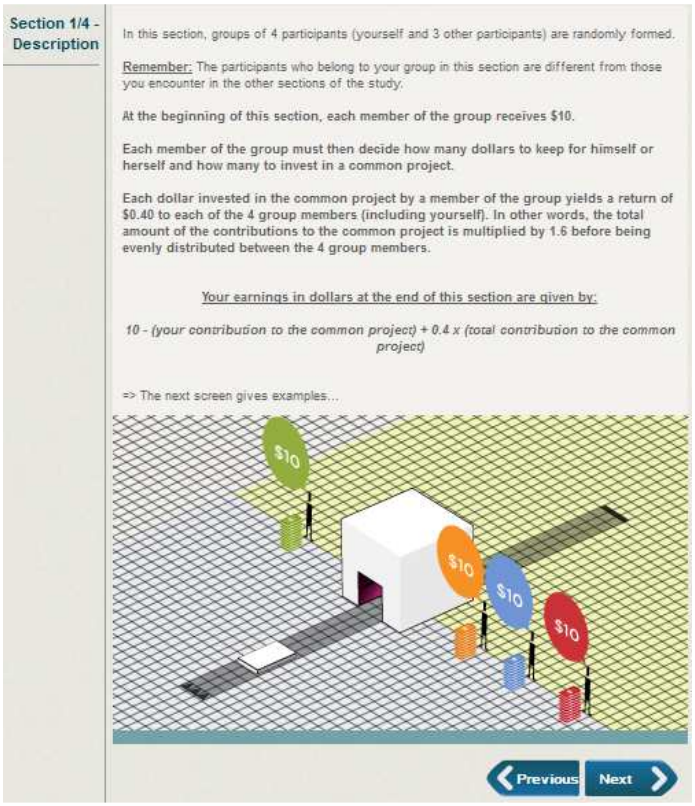

Figure 2. A typical Barnstar

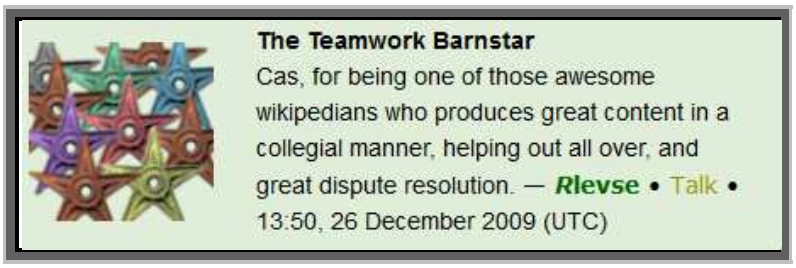

Figure 3. The decision screen of the conditional Public Goods game

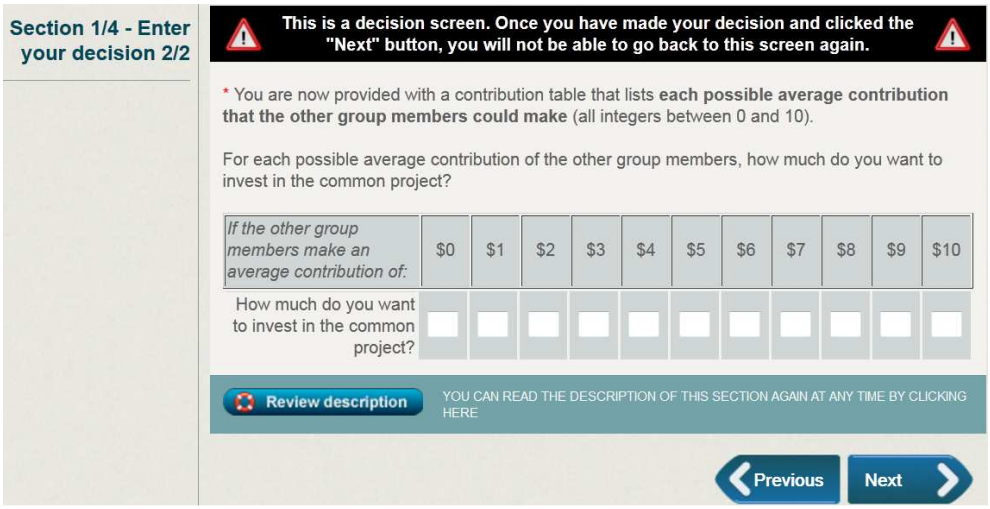

Figure 4. The Wikipedia recruitment banner

Figure 5. Distribution of the number of Wikipedia contributions per experimental group: whole population vs. study participants

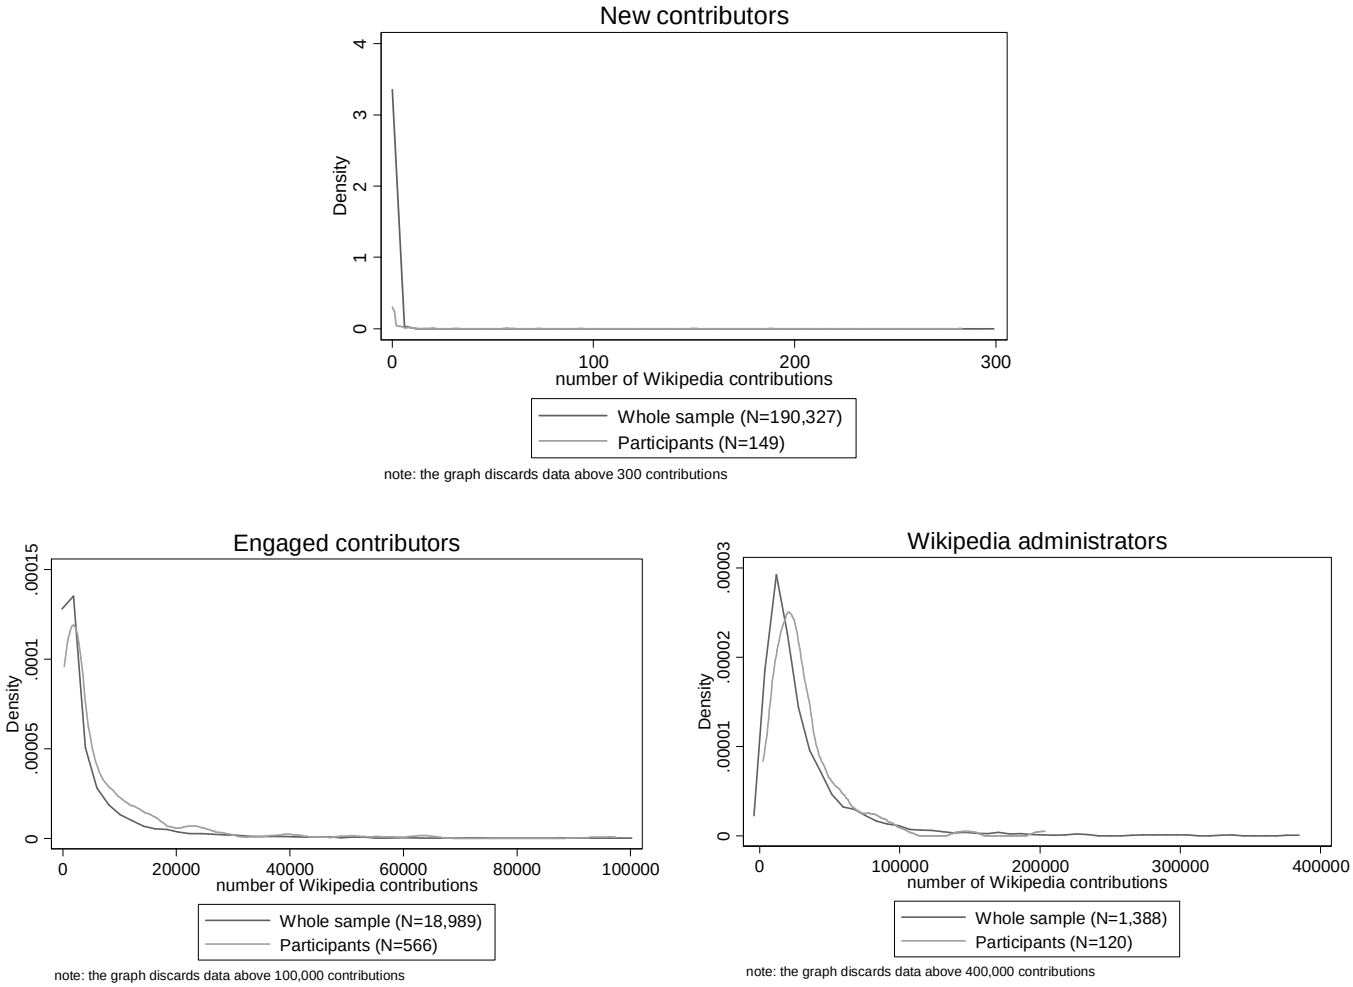

Notes: Kernel density estimates.

Table 1. Descriptive statistics

	New contributors	Engaged contributors	Administrators	Other
Number of observations (*N*)	149	566	120	15
DEPENDENT VARIABLE				
Mean – number of Wikipedia contributions	8.64	9719.83	41229.24	543.13
	(3.56)	(23519.39)	(86191.33)	(1664.19)
	[0; 273]	[303; 364157]	[2475; 922895]	[0; 6547]
SOCIAL PREFERENCES MEASURES				
Altruism (N=405)				
(i) Proportion of endowment transferred – Dictator	0.38	0.36	0.42	0.28
	(0.30)	(0.30)	(0.28)	(0.26)
(ii) Proportion of endowment transferred – in-group Dictator	0.46	0.45	0.48	0.40
	(0.29)	(0.30)	(0.28)	(0.20)
Reciprocity (N=850 & N=445)				
(i) Average proportion of endowment conditionally contributed – Public Goods	0.45	0.54	0.52	0.52
	(0.27)	(0.24)	(0.24)	(0.26)
(ii) Average proportion of amount returned – Trust	0.45	0.51	0.46	0.54
	(0.25)	(0.23)	(0.23)	(0.22)
Social image (N=850 & N=456)				
(i) Mean size of Wikipedia user page (in bytes)	453.81	5586.24	9179.64	1238.60
	(3597.62)	(10859.97)	(11012.09)	(3438.12)
Number of Barnstars receivers	4	340	109	3
Mean - number of Barnstars received	1.5	6.14	16.8	5
	(0.58)	(8.57)	(15.99)	(6.93)
(ii) Proportion signaling Barnstars	0.25	0.50	0.70	0.33
	(0.50)	(0.50)	(0.46)	(0.58)
DEMOGRAPHIC VARIABLES				
Age	27	34	34	33
	(11.81)	(14.73)	(12.86)	(8.84)
Proportion female	0.15	0.09	0.11	0.07
	(0.36)	(0.29)	(0.31)	(0.26)
Degree level	3.97	4.55	4.88	4.73
	(1.92)	(1.80)	(1.64)	(1.75)
Salary level	3.17	3.80	4.01	3.79
	(2.15)	(2.34)	(2.25)	(2.12)
Risk aversion level	6.16	5.66	5.53	4.67
	(2.36)	(2.34)	(2.38)	(2.09)

Notes: Standard errors are reported in parenthesis. *mean number of Wikipedia contributions* = mean number of modifications implemented in Wikipedia (minimum and maximum values are reported in brackets). *Degree level*: 1 = "less than high school"; 2 = "high school"; 3 = "some college"; 4 = "2 years college degree"; 5 = "4 years college degree (BA, BS)"; 6 = "masters degree"; 7 = "professional degree (MD, JD)"; 8 = "doctoral degree". *Salary level* (monthly): 1 = "0 USD"; 2 = "less than 1000 USD"; 3 = "between 1000 and 2000 USD"; 4 = "between 2000 and 3000 USD"; 5 = "between 3000 and 4000 USD"; 6 = "between 4000 and 5000 USD"; 7 = "between 5000 and 7500 USD"; 8 = "between 7500 and 10000 USD"; 9 = "more than 10000 USD". *Risk aversion level* = whether subjects generally see themselves as fully prepared to take risks as opposed to generally trying to avoid taking risks: 0 = "unwilling to take risks" to 10 = "fully prepared to take risks".

Table 2. Sample common demographic characteristics: Wikimedia editor survey vs. our study

	2011 Wikimedia editor survey	Our study
Age		
12 to 17	13%	4%
18 to 21	14%	17%
22 to 29	26%	30%
30 to 39	19%	20%
40 or more	28%	28%
Gender		
Proportion female	9%	10%
Education level		
Primary	9%	5%
Secondary	30%	31%
Bachelors / associate	35%	34%
Master's	18%	22%
PhD	8%	7%

Notes: The Wikimedia editor survey excludes respondents under 12 and over 82 from the sample. The age and gender statistics are based on the population of respondents with a positive number of Wikipedia contributions (N=4,930). The Education level statistics are based on the whole population of respondents (N=5,073). In this table, we base our own statistics on the same calculation rules.

Table 3. Number of Wikipedia contributions and demographic characteristics

	(1)	(2)	(3)	(4)
Dependent variable: number of Wikipedia contributions	Whole sample	Below median	Above median	Admins
age	0.0167***	0.0136***	0.00516	0.0160***
	(0.00306)	(0.00518)	(0.00338)	(0.00541)
female	-0.365**	-0.665***	-0.0332	-0.244
	(0.147)	(0.243)	(0.152)	(0.249)
degree level	0.0582**	0.0971***	0.0165	-0.0391
	(0.0246)	(0.0362)	(0.0309)	(0.0446)
salary level	0.00282	-0.00846	-0.0133	-0.0509*
	(0.0200)	(0.0313)	(0.0225)	(0.0305)
Risk aversion	-0.0325*	-0.0640**	-0.00170	0.00533
	(0.0169)	(0.0259)	(0.0193)	(0.0279)
Constant	8.310***	5.901***	9.420***	10.46***
	(0.175)	(0.266)	(0.189)	(0.291)
N	649	325	324	113
Pseudo R^2	0.00507	0.00698	0.000488	0.00385

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Model (1) is all non admin subjects; model (2) is non admin subjects below the median number of Wikipedia contributions (i.e. 1905 contributions); model (3) is non admin subjects above the median number of Wikipedia contributions; model (4) is all admin subjects.

Table 4. Number of Wikipedia contributions and altruism motive

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable: number of Wikipedia contributions	Whole sample	Whole sample	Below median	Below median	Above median	Above median	Admins	Admins
Altruism (Dictator)	-0.183		-0.184		-0.239		0.181	
	(0.207)		(0.332)		(0.224)		(0.317)	
Altruism (Dictator in-group)		-0.184		-0.177		-0.319		0.199
		(0.207)		(0.332)		(0.222)		(0.353)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	305	305	159	159	146	146	56	56
Pseudo R^2	0.00698	0.00699	0.00834	0.00833	0.00226	0.00257	0.0112	0.0112

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Altruism (Dictator)* = proportion of endowment transferred in the Dictator game. *Altruism (Dictator in-group)* = proportion of endowment transferred in the directed Dictator game. Models (1) and (2) are all non admin subjects; models (3) and (4) are non admin subjects below the median number of Wikipedia contributions (i.e. 1905 contributions); models (5) and (6) are non admin subjects above the median number of Wikipedia contributions; models (7) and (8) are all admin subjects.

Table 5. Number of Wikipedia contributions and reciprocity motive

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable: number of Wikipedia contributions	Whole sample	Whole sample	Below median	Below median	Above median	Above median	Admins	Admins
Reciprocity (Public Goods)	0.378**		0.796***		-0.107		-0.631**	
	(0.162)		(0.246)		(0.187)		(0.307)	
Reciprocity (Trust)		0.443*		1.136***		0.0424		-0.990**
		(0.242)		(0.392)		(0.273)		(0.447)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	649	344	325	166	324	178	113	57
Pseudo R^2	0.00554	0.00516	0.00959	0.0142	0.000535	0.000572	0.00538	0.00594

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Reciprocity (Public Goods)* = average proportion of endowment conditionally contributed in the Public Goods game strategy method; *Reciprocity (Trust)* = average proportion of amount received that is returned by the subject in the Trust game strategy method. Models (1) and (2) are all non admin subjects; models (3) and (4) are non admin subjects below the median number of Wikipedia contributions (i.e. 1905 contributions); models (5) and (6) are non admin subjects above the median number of Wikipedia contributions; models (7) and (8) are all admin subjects.

Table 6. Number of Wikipedia contributions and social image motive

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: number of Wikipedia contributions	Whole sample	Below median	Above median	Whole sample	Admins	Admins
Social signaler (user page)	1.305***	1.805***	0.261***		0.354***	
	(0.0845)	(0.126)	(0.101)		(0.135)	
Social signaler (Barnstars)				0.288***		0.089
				(0.0969)		(0.151)
nb Barnstars				0.0405***		0. 0319***
				(0.00287)		(0.00434)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
N	649	325	324	308	113	102
Pseudo R^2	0.0256	0.0467	0.00150	0.0192	0.00640	0.0358

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Social signaler (Barnstars)* = 1 if the subject decided to advertise at least one of his Barnstars on his user page (0 otherwise). *Social signaler (user page)* = 1 if the subject has a Wikipedia user page whose size (in bytes) is greater than the median in the sample of all non admin subjects. *nb Barnstars* = total number of Barnstars received by each subject. Model (1) is all non admin subjects; models (2) is non admin subjects below the median number of Wikipedia contributions (i.e. 1905 contributions); models (3) is non admin subjects above the median number of Wikipedia contributions; model (4) is all non admin subjects who received Barnstars; models (5) and (6) are all admin subjects.

Table 7. Interaction between reciprocity and social image motives – regular contributors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable: number of Wikipedia contributions	Whole sample	Whole sample	Below median	Below median	Above median	Above median	Whole sample	Whole sample
Social signaler (user page)	1.648***	2.021***	2.096***	2.133***	0.678***	0.562		
	(0.190)	(0.281)	(0.270)	(0.466)	(0.230)	(0.348)		
Social signaler (Barnstars)							0.665***	0.812**
							(0.224)	(0.335)
Reciprocity (Public Goods) x Social signaler (user page)	-0.0987		0.354		-0.442*			
	(0.198)		(0.304)		(0.228)			
Reciprocity (Public Goods) x non Social signaler (user page)	0.563**		0.932***		0.338			
	(0.250)		(0.318)		(0.321)			
Reciprocity (Trust) x Social signaler (user page)		-0.294		0.632		-0.0176		
		(0.309)		(0.586)		(0.316)		
Reciprocity (Trust) x non Social signaler (user page)		1.067***		1.173**		0.451		
		(0.376)		(0.501)		(0.507)		
Reciprocity (Public Goods) x Social signaler (Barnstars)							-0.324	
							(0.242)	
Reciprocity (Public Goods) x non Social signaler (Barnstars)							0.349	
							(0.263)	
Reciprocity (Trust) x Social signaler (Barnstars)								-0.614*
								(0.337)
Reciprocity (Trust) x non Social signaler (Barnstars)								0.154
								(0.497)
nb Barnstars							0.0400***	0.0371***
							(0.00288)	(0.00380)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	649	344	325	166	324	178	308	164
Pseudo R^2	0.0261	0.0270	0.0491	0.0562	0.00221	0.00214	0.0197	0.0223

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Social signaler (Barnstars)* = 1 if the subject decided to advertise at least one of his Barnstars on his user page (0 otherwise). *Social signaler (user page)* = 1 if the subject has a Wikipedia user page whose size (in bytes) is greater than the median in the sample of all non admin subjects. *Reciprocity (Public Goods)* = average proportion of endowment conditionally contributed in the Public Goods game strategy method; *Reciprocity (Trust)* = average proportion of amount received that is returned by the subject in the Trust game strategy method. *nb Barnstars* = total number of Barnstars received by each subject. Models (1) and (2) are all non admin subjects; models (3) and (4) are non admin subjects below the median number of Wikipedia contributions (i.e. 1905 contributions); models (5) and (6) are non admin subjects above the median number of Wikipedia contributions; models (7) and (8) are all non admin subjects who received Barnstars.

Table 8. Interaction between reciprocity and social image motives – Wikipedia administrators

	(1)	(2)	(3)	(4)
Dependent variable: number of Wikipedia contributions	Admins	Admins	Admins	Admins
Social signaler (user page)	0.601*	0.364		
	(0.324)	(0.485)		
Social signaler (Barnstars)			-0.245	0.0149
			(0.296)	(0.464)
Reciprocity (Public Goods) x Social signaler (user page)	-0.854**			
	(0.389)			
Reciprocity (Public Goods) x non Social signaler (user page)	-0.370			
	(0.441)			
Reciprocity (Trust) x Social signaler (user page)		-1.070**		
		(0.546)		
Reciprocity (Trust) x non Social signaler (user page)		-0.686		
		(0.788)		
Reciprocity (Public Goods) x Social signaler (Barnstars)			-1.041***	
			(0.328)	
Reciprocity (Public Goods) x non Social signaler (Barnstars)			-1.293***	
			(0.461)	
Reciprocity (Trust) x Social signaler (Barnstars)				-0.981**
				(0.458)
Reciprocity (Trust) x non Social signaler (Barnstars)				-0.159
				(0.978)
nb Barnstars			0.0195***	0.0339***
			(0.00342)	(0.00716)
Constant	10.55***	10.05***	10.35***	10.77***
	(0.329)	(0.596)	(0.340)	(0.654)
Control variables	Yes	Yes	Yes	Yes
N	102	49	113	57
Pseudo R^2	0.0172	0.0201	0.00832	0.00670

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Social signaler (Barnstars)* = 1 if the subject decided to advertise at least one of his Barnstars on his user page (0 otherwise). *Social signaler (user page)* = 1 if the subject has a Wikipedia user page whose size (in bytes) is greater than the median in the sample of all non admin subjects. *Reciprocity (Public Goods)* = average proportion of endowment conditionally contributed in the Public Goods game strategy method; *Reciprocity (Trust)* = average proportion of amount received that is returned by the subject in the Trust game strategy method. *nb Barnstars* = total number of Barnstars received by each subject. Models (1) to (4) are all admin subjects.

Table 9. Patterns of Wikipedia contributions and generalized trust

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable :	number of Wikipedia contributions	number of Wikipedia contributions	number of Wikipedia contributions	number of Wikipedia contributions	number of users blocked	number of pages deleted	number of pages protected	Time spent on admin activities
	Whole sample	Below median	Above median	Admins	Admins	Admins	Admins	Admins
Trust	0.0780	-0.0265	-0.0393	-0.730**	-1.004**	-0.725*	-0.626	-3.703*
	(0.180)	(0.272)	(0.187)	(0.309)	(0.463)	(0.419)	(0.460)	(1.818)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	305	159	146	56	56	56	56	27
Pseudo / Adj. R^2	0.00687	0.00819	0.00191	0.0150	0.00791	0.00516	0.00830	0.247

Notes: Negative binomial estimates (except for column (8), which features OLS). Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Trust* = proportion of endowment sent in the Trust game. Column (1) is all non admin subjects; column (2) is non admin subjects below the median number of Wikipedia contributions (i.e. 1905 contributions); column (3) is non admin subjects above the median number of Wikipedia contributions; columns (4), (5), (6), (7) and (8) is all admin subjects. *Time spent on admin activities* = answer to the question: "what fraction of the time that you spend working on Wikipedia do you specifically devote to activities that admins only can perform (e.g. deleting and protecting pages, blocking and unblocking users etc.) as opposed to the regular editing activities mentioned above? Please choose one number on the following scale, where 0 means "I do not spend any of my working time on Wikipedia performing admin-related tasks" and 10 means "I spend all of my working time on Wikipedia performing admin-related tasks"."