



HAL
open science

Online cooperation and peer production

Jérôme Hergueux

► **To cite this version:**

Jérôme Hergueux. Online cooperation and peer production. Economics and Finance. Université de Strasbourg, 2014. English. NNT : 2014STRAB003 . tel-01133392

HAL Id: tel-01133392

<https://theses.hal.science/tel-01133392>

Submitted on 19 Mar 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



UNIVERSITY OF STRASBOURG

AUGUSTIN COURNOT DOCTORAL SCHOOL

LaRGE

Ph.D. DISSERTATION

Jérôme Hergueux

Defended : 12 / 06 / 2014

To get the title of : **Doctor of the University of Strasbourg**

Discipline : Economics

Online Cooperation and Peer Production

ADVISORS:

Yann Algan

Professor of Economics, Sciences Po

Laurent Weill

Professor of Economics, University of Strasbourg

REFEREES:

Guillaume Fréchette

Professor of Economics, New-York University

Marie Claire Villeval

Research Professor in Economics, CNRS

OTHER COMMITTEE MEMBERS:

Yochai Benkler

Berkman Professor of Entrepreneurial Studies, Harvard University

Paul Seabright

Professor of Economics, Toulouse School of Economics

Online Cooperation and Peer Production

Abstract

Internet is a very attractive technology for the implementation of experiments. It allows to obtain larger and more diverse samples and gives the researcher the opportunity to extract from the Internet a wealth of field data that document the real-world decisions and behavior of his subjects. Notwithstanding those appealing features, the development of the “online laboratory” remains in its infancy, mainly because of the threats to validity and important practical challenges typically associated with online experimentation.

More than a potentially powerful medium to run experiments, the Internet is also a very promising field of economic research. Over the past 20 years, its diffusion has significantly reduced communication costs and increased information flows between economic agents. This technological change has notably fostered the emergence of a new production model – peer production – which is primarily based on voluntary contributions and large-scale collaboration. Peer production is a significant organizational innovation: agents voluntarily self-assign work and successfully coordinate towards the provision of global public goods, in the absence of price signals and without any pre-specified design rule or formal leadership.

From Open Source Software to Wikipedia, peer production involves hundreds of thousands of contributors worldwide. It is an important source of value creation in the most competitive sectors of information and technology, as well as a major source of innovation. Beyond its economic significance, the emergence of peer production also represents an opportunity to shed new lights on a number of longstanding but notably difficult questions in the literature. Given the unconventional nature of many of the work incentives at play in peer production environments, those are particularly well suited for researching the impact of non standard economic preferences on public goods provision, studying their role as work incentives, and assessing their consequences in terms of organizational economics.

The first contribution of this dissertation is a methodological one. Chapter 1 develops and assesses the reliability of a novel online experimentation tool specifically designed to strengthen the internal validity of the decisions elicited over the Internet. Chapter 2 and 3 document the rise of peer production as a new and significant model for organizing production. Exploiting the context of peer production, those chapters leverage the online experimentation tool developed in Chapter 1 and rely on a combination of large-scale online experiments and computational methods (i.e. the systematic extraction of data on subjects’ field behavior) to respectively (i) provide the first comprehensive field test of the theory of the private provision of public goods, (ii) study the importance of social preferences as work motives within real-world productive organizations and (iii) report the first field evidence of endogenous sorting behavior of economic agents within productive teams based on their cooperative types.

Keywords: Field Experiment, Social Preferences, Public Goods, Labor Economics, Peer Production, Wikipedia, Open Source Software, Internet, Methodology

Résumé

Internet est une technologie très attractive pour la mise en place d'expériences. Il permet d'obtenir des échantillons plus grands et plus divers, et donne au chercheur l'opportunité d'extraire d'Internet toute une série de données de terrain qui documentent les décisions et le comportement de ses sujets. Malgré ces caractéristiques attrayantes, le développement du "laboratoire en ligne" en reste à ses balbutiements, principalement du fait des menaces à la validité et des importantes difficultés pratiques liées à l'expérimentation en ligne.

Plus qu'un outil potentiellement puissant pour la mise en place d'expériences, Internet est aussi un terrain de recherche économique très prometteur. Durant les 20 dernières années, sa diffusion a significativement réduit les coûts de communication et augmenté les échanges d'information entre agents économiques. Cette évolution technologique a favorisé l'émergence d'un nouveau modèle de production – la production par les pairs – basée prioritairement sur les contributions volontaires et la collaboration à large échelle. La production par les pairs est une innovation organisationnelle significative: les agents s'auto-assignent des tâches et se coordonnent avec succès vers la production de biens publics globaux, en l'absence de signaux de prix et sans règle de conception préétablie ou leadership formel.

Des logiciels Open Source à Wikipédia, la production par les pairs mobilise des centaines de milliers de contributeurs de par le monde. C'est une source importante de création de valeur dans les secteurs très compétitifs de l'information et de la technologie, ainsi qu'une source majeure d'innovation. Au-delà même de son importance économique, l'émergence de la production par les pairs représente une opportunité d'éclairer un certain nombre de questions anciennes et particulièrement ardues dans la littérature d'un jour nouveau. Compte-tenu de la nature souvent non conventionnelle des incitations au travail dans les environnements de production par les pairs, ceux-ci sont particulièrement adaptés à l'étude de l'impact des préférences économiques non standard sur la production de biens publics, à l'analyse de leur rôle en tant que motivations au travail, ainsi qu'à l'évaluation de leurs conséquences en termes d'économie organisationnelle.

La première contribution de ce travail de thèse est d'ordre méthodologique. Le chapitre 1 développe et évalue la fiabilité d'un nouvel outil d'expérimentation en ligne, construit spécifiquement de manière à renforcer la validité interne des décisions élicitées sur Internet. Les chapitres 2 et 3 documentent l'émergence de la production par les pairs en tant qu'un modèle nouveau et significatif d'organisation de la production. Exploitant le contexte de la production par les pairs, ces chapitres utilisent l'outil d'expérimentation en ligne développé dans le chapitre 1 et s'appuient sur une combinaison d'expériences en ligne à large échelle et de méthodes computationnelles (i.e. l'extraction systématique de données sur le comportement de terrain des sujets) afin de (i) mener le tout premier test de terrain exhaustif de la théorie de la production privée de biens publics, (ii) étudier l'importance des préférences sociales en tant que motivations au travail au sein d'organisations productives réelles et (iii) procéder aux premiers tests de terrain documentant des comportements endogènes d'appariement des agents économiques au sein d'équipes productives en fonction de leur type coopératif.

Mots-clés: Expérience de terrain, Préférences Sociales, Biens Publics, Economie du Travail, Production par les Pairs, Wikipédia, Logiciels Open Source, Internet, Méthodologie

Word of thanks

I first would like to thank my advisors, Yann Algan and Laurent Weill, who both gave me their trust as well as a great deal of freedom and autonomy while pursuing this research project. It helped me grow, and I needed it to complete this research endeavor.

I would like to thank Yochai Benkler, who has inspired me during all those years, academically and well beyond.

I thank Guillaume Fréchette and Marie Claire Villeval for accepting to review my work, as well as Paul Seabright for accepting to preside over my dissertation committee. It is an honor for me to have such distinguished researchers as committee members.

I thank Nicolas Jacquemet and Emeric Henry, who have both greatly contributed to my research efforts.

* * * *

I would like to thank my Mom and my Dad, who, both in their own way, allowed me to be where I am and become the person that I am today.

I would like to thank Sibylle, the ocean of gentleness that helped me go through the most difficult moments.

I thank my friends from Elsass. They will recognize each other and know how important they are to me.

I thank my Iraqi – French family. That's it! Al Doctora!

* * * *

At Harvard:

I would like to thank my roommates and closest Harvard friends: Mitra Akhtari, Gabriel Smagghue and Jack Willis. It really wouldn't have been the same without you.

I thank the whole Berkman staff and Geeks for their cheerfulness, their support and their help throughout my work. Onward!

I thank the wonderful community of Fellows with which I've been fortunate to interact and exchange so many fascinating ideas during those past 3 years.

I thank the Berkman Cooperation group, especially my friends: Mayo Fuster-Morell, Benjamin Mako Hill, Aaron Shaw and Dennis Yi Tenen. You cannot underestimate your contribution to this work.

At Sciences Po:

I thank all the people with whom I've been fortunate to work and interact, especially at the médialab. In particular: Barbara Bender, Paul Girard, Romain Guillebert, Andrews-Junior Kimbembe and Matt Van Antwerp.

I thank my officemates and friends: Elizabeth Beasley, Thomas Grjebine, Anne l'Hôte and Matthieu Solignac.

At the University of Strasbourg:

I thank the whole Faculty at the Institute for Political Studies. You were my first Professors. In a sense, I am proud to become your colleague.

I thank my officemates and friends: Anais Hamelin, Pierre Pessarossi and Blandine Zimmer.

Finally, I thank the vibrant Wikimedia and Open Source Software communities. You inspired me, challenged me, got me thinking and, from time to time, made my life miserable...

Remerciements

J'aimerais tout d'abord remercier mes directeurs de thèse, Yann Algan and Laurent Weill, qui m'ont tout deux accordé leur confiance ainsi qu'une très grande liberté et autonomie dans la conduite de ce projet de recherche. Cela m'a aidé à grandir, et j'en avais besoin pour réussir dans cette entreprise.

J'aimerais remercier Yochai Benkler, qui m'a inspiré pendant toutes ces années, académiquement et bien au delà.

Je remercie Guillaume Fréchette et Marie Claire Villeval d'avoir accepté de commenter mon travail, ainsi que Paul Seabright d'avoir accepté de présider ce jury de thèse. C'est un honneur d'avoir d'aussi remarquables chercheurs comme membres de jury.

Je remercie Nicolas Jacquemet et Emeric Henry, qui ont tous deux beaucoup contribué à mes efforts de recherche.

* * * *

J'aimerais remercier ma Maman et mon Papa, qui, chacun à leur manière, m'ont permis d'être où je suis et de devenir la personne que je suis aujourd'hui.

J'aimerais remercier Sibylle, l'océan de douceur qui m'a aidé à traverser les moments les plus difficiles.

Je remercie mes copains d'Elsass. Ils se reconnaîtront et savent à quel point ils comptent pour moi.

Je remercie ma famille Franco-Irakienne. Ca y est! Al Doctora!

* * * *

A Harvard:

J'aimerais remercier mes colocataires et amis d'Harvard les plus proches: Mitra Akhtari, Gabriel Smagghue et Jack Willis. Rien n'aurait été pareil sans vous.

Je remercie l'ensemble du staff et les Geeks de Berkman pour leur bonne humeur, leur soutien et leur aide tout au long de mon travail. Onward!

Je remercie la fantastique communauté de Fellows avec laquelle j'ai eu la chance d'interagir et d'échanger tant d'idées fascinantes pendant ces trois dernières années.

Je remercie le Berkman Cooperation group, tout particulièrement mes amis: Mayo Fuster-Morell, Benjamin Mako Hill, Aaron Shaw, et Dennis Yi Tenen. Vous ne pouvez sous-estimer votre contribution à ce travail.

A Sciences Po:

Je remercie tous ceux avec qui j'ai eu la chance de travailler et d'interagir, particulièrement au médialab. En particulier: Barbara Bender, Paul Girard, Romain Guillebert, Andrews-Junior Kimbembe et Matt Van Antwerp.

Je remercie mes collègues de bureau et amis: Elizabeth Beasley, Thomas Grjebine, Anne l'Hôte et Matthieu Solignac.

A l'Université de Strasbourg:

Je remercie l'ensemble du corps enseignant de l'Institut d'Etudes Politiques. Vous avez été mes premiers professeurs. En un sens, je suis fier de devenir votre collègue.

Je remercie mes collègues de bureau et amis: Anais Hamelin, Pierre Pessarossi et Blandine Zimmer.

Finalement, je remercie les vibrantes communautés Wikimedia et du Logiciel Open Source. Vous m'avez inspiré, remis en question, fait réfléchir, et, parfois, m'avez rendu la vie impossible...

Notice

The three chapters of this dissertation are self-contained research articles. This explains that the term "paper" is used, and why some information can be redundant. The first chapter of this dissertation has been accepted for publication and is forthcoming in *Experimental Economics*. The second and third chapters are work in progress. The second chapter will soon be submitted to an academic journal.

Avertissement

Les trois chapitres de cette thèse sont des articles de recherche indépendants. Ceci explique la présence du terme "paper", ainsi que l'éventuelle redondance de certaines informations. Le premier chapitre de cette thèse a été accepté pour publication et est à paraître dans *Experimental Economics*. Les deuxième et troisième chapitres sont un travail en cours. Le second chapitre sera bientôt soumis à un journal académique.

Contents

Abstract / Résumé	iii
Word of thanks / Remerciements	vi
Notice / Avertissement	xi
Contents	xiii

General introduction & research statement.....	1-9
--	-----

CHAPTER 1 – SOCIAL PREFERENCES IN THE ONLINE LABORATORY: A RANDOMIZED EXPERIMENT

Abstract.....	12
1 Introduction.....	13
2 Design of the experiment.....	14
2.1 The decision problems.....	15
2.2 Procedures common to both implementations.....	20
2.3 Practical implementation of the experiment.....	19
2.4 Summary of the design.....	20
3 Social preferences in the online laboratory.....	21
4 Do subjects actually behave more pro-socially online?.....	23
4.1 Internal validity of the comparison.....	23
4.2 Differences in underlying secondary outcomes.....	25
4.3 The effect of Internet-specific differences in design.....	26
4.3.1 Design of the robustness treatments.....	27
4.3.2 Results.....	27
5 Discussion.....	29
6 Conclusion.....	31
References.....	33
Tables and figures.....	35
Appendix.....	48

CHAPTER 2 – COOPERATION IN A PEER PRODUCTION ECONOMY: EXPERIMENTAL EVIDENCE FROM WIKIPEDIA

Abstract.....	51
1 Introduction.....	52
2 Background on Wikipedia.....	53
3 Design of the study.....	54
3.1 Measuring social preferences.....	58

3.2	Experimental procedures.....	60
3.3	Implementation of the experiment.....	61
4	Results.....	63
4.1	Descriptive statistics and analysis.....	63
4.2	Prosocial foundations of cooperation: regular Wikipedia contributors.....	65
4.3	Prosocial foundations of cooperation: Wikipedia administrators.....	67
5	Discussion.....	69
6	Conclusion.....	71
	References.....	73
	Tables and figures.....	77
	Appendix.....	87

**CHAPTER 3 – SOCIAL MOTIVES AND THE ORGANIZATION OF PRODUCTION:
EXPERIMENTAL EVIDENCE FROM OPEN SOURCE SOFTWARE**

	Abstract.....	90
1	Introduction.....	91
2	Design of the study.....	95
2.1	Measuring motivations.....	95
2.1.1	Experimental measures.....	95
2.1.2	Survey measures.....	97
2.2	Experimental procedures.....	98
2.3	Implementation of the experiment.....	100
2.4	Dependent variables and identification strategy.....	101
3	Developers’ motivations to contribute to open source.....	104
3.1	Descriptive statistics and analysis.....	104
3.2	Motivations to contribute to OSS.....	106
3.3	Intensive and extensive margins of contributions.....	108
4	Developers’ sorting behavior at the team level.....	109
5	Conclusion.....	111
	References.....	113
	Tables and figures.....	117
	Appendix.....	126

	Concluding remarks and future work.....	129-130
--	---	---------

General introduction & research statement

We live an increasing part of our lives over the Internet. From discussion forums and social networking sites (e.g. Facebook) to massively multiplayer online role-playing (e.g. World of Warcraft), from online auction sites (e.g. eBay) to online labor markets (e.g. ODesk), the cyberspace is now a prominent part of the “real” social and economic world.

In the field of experimental economics, it has been a long time since researchers have called upon the development of the “online laboratory”. There are three main reasons for experimentalists’ interest in online experimentation.

First, the Internet is appealing because it allows to reach more diverse samples, recruit larger and sometimes more representative subject pools and even conduct cross-cultural social experiments in real time at an affordable cost. In a seminal paper, Henrich et al. (2010) warn against behavioral scientists’ current over-reliance on data overwhelmingly gathered from populations of undergraduate students and recommend a major effort in broadening the sample base:

“Although we are certainly not the first to worry about the representativeness of prevalent undergraduate samples in the behavioral sciences, our efforts to compile an empirical case have revealed an even more alarming situation than previously recognized. The sample of contemporary Western undergraduates that so overwhelms our database is not just an extraordinary restricted sample of humanity; it is frequently a distinct outlier vis-à-vis other global samples” (p. 82).

Table 1 documents the number of Internet users, Internet penetration rate and the growth rate of the population of Internet users by regions of the world as of June 2012. The picture seems to make a compelling case for the use of the Internet as a *medium* to conduct experiments with large and diverse samples: it is now possible to reach 78.6% of the North American population through the Internet with relative ease, and while only 15.6% of the African population can currently be reached through this method, the exponential growth of its user base will soon make it an attractive tool for conducting experiments in the developing world.

Second, as we spend a significant fraction of our time in the network, the Internet becomes a prominent experimental *field* of research in its own right (see Bainbridge 2007). It therefore makes a lot of sense to conduct experiments directly over the Internet to understand the various types of social and economic activities that people engage in online. An increasing number of economics papers make use of field experiments (Levitt & List 2009). Those are appealing because, contrary to laboratory experiments, they “creat[e] a context that is similar to the one in which economic agents operate” (Loewenstein 1999). As a result, field experiments are thought to be more externally valid, in the sense that their results generalize more easily to the “real world” phenomena from which researchers seek to learn. Technically speaking, the attractiveness of the Internet to conduct field economic experiments is further reinforced by the fact that it appears as a natural environment for subjects to interact with varying degrees of anonymity and, if necessary, without repetition.

Last, conducting experiment over the Internet allows the researcher to *relate his experimental results to detailed real-world records of his subjects’ field behavior*, therefore helping bridging the gap between experimental and observational data. Indeed, in many instances, it is possible to extract from the

Internet a wealth of externally valid information on individual's decisions and trace their behavior. Depending on the assumptions made or on the research question asked by the researcher, this feature of the Internet can allow to either relieve the traditional tension between internal and external validity in economic experiments, or to assess the ecological validity of experimental procedures that are used extensively in the lab to test economic theory. This feature of the Internet makes it a very promising field of research for economists and other social scientists. In a seminal paper, Lazer et al. (2009) call for some massive investments in the systematic collection and analysis of the digital world's vastly untapped data for the social sciences:

"The capacity to collect and analyze massive amounts of data has unambiguously transformed such fields as biology and physics. The emergence of such a data-driven "computational social science" has been much slower, largely spearheaded by a few intrepid computer scientists, physicists, and social scientists. If one were to look at the leading disciplinary journals in economics, sociology, and political science, there would be minimal evidence of an emerging computational social science engaged in quantitative modeling of these new kinds of digital traces." (p.2)

Notwithstanding those appealing features, however, and although the Internet is now a well-established technology in the developed world, the development of the "online laboratory" still remains in its infancy. To be sure, there are some important methodological and practical challenges associated with online experimentation. The first contribution of this dissertation will therefore be a methodological one. It will report on the construction and evaluation of a novel tool to conduct online experiments.

One of the most interesting aspects of the Internet is that the diffusion and interconnection of personal computers has significantly reduced communication costs and increased information flows between economic agents. This technological evolution has led to an increase in market efficiency and, arguably, social welfare. ODesk – an online market for skilled labor in which computer programmers can bid to execute modular contracts posted by customers – is a good example of the increase in market efficiency that a significant reduction in communication costs can foster.

Beyond market efficiency gains, however, the Internet has also fostered the emergence of a qualitatively distinct model of production – peer production – alongside firms, markets and the State, which is primarily based on voluntary contributions and large-scale collaboration (Benkler 2002, 2006). Peer production represents a significant organizational innovation in which agents voluntarily self-assign work and successfully coordinate towards the provision of global public goods, in the absence of price signals and without any pre-specified design rule or formal leadership. During the past two decades this emerging organizational model has been an important source of value creation in the most innovative sectors of information and technology.

Peer production is therefore worth studying in and of itself. Wikipedia, the first experimental field of this dissertation, currently hosts over 25 million freely usable articles in 285 languages. Its revealed informational value seems to be enormous to society. It is the 5th most visited website on the Internet, receiving over 500 million unique visitors per month worldwide. 60% of European doctors declare using Wikipedia for professional purposes, and an early evaluation of the quality of its scientific

entries actually found them to be practically indistinguishable from those in the professionally edited encyclopedia Britannica (Giles 2005).

Open Source Software development platforms, the second experimental field of this dissertation, involve an estimated 800,000 developers around the world. Open Source Software (OSS) is a major source of innovation, which is responsible for most of the basic utilities on which the Internet runs (e.g. the Apache web server, Sendmail, the Domain Name System management software BIND), popular programming languages (e.g. Python, Perl) and programming environments (e.g. Eclipse). OSS also successfully competes with many of its firm-based counterparts in the realm of enterprise systems (e.g. Linux, which was adopted by a number of large commercial firms such as IBM, Apple and Sun for its relative reliability, resilience to virus attacks, and bug correcting speed) and end-user applications (e.g. Android, OpenOffice, VLC Media Player, Mozilla Firefox). The economic value generated by OSS is estimated to be substantial. Indeed, Walli, Gynn, and Rotz (2005) report that 87% of US businesses now rely on OSS for some of their daily activities. Ghosh (2007) estimates the cost of recreating the existing open source code at 12 billion euros, while Greenstein and Nagle (2014) estimate the value of the Apache web server alone to range between 2 and 12 billion dollars.

But even beyond its relative novelty and economic significance, peer production also represents an opportunity to shed new lights on a number of prominent but notably difficult questions in the literature. As Benkler (2013) notes:

“The implications of peer production are broader than the direct economic impact of the practice. Beyond the magnitude of its effects on innovation and knowledge production in the networked economy and participation in the networked society, the success of peer production and online cooperation has several implications for economics more generally. It requires that we refine our ideas about motivation or incentives; it recalibrates the roles of property and contract [...] in the growth-critical domains of knowledge-dependent production and innovation; and it requires adaptations to the theory of the firm.” (p. 2)

Given the unconventional nature of many of the work incentives at play in peer production environments, this dissertation exploits the context of peer production to research the impact of non standard economic preferences on public goods provision, study their role as work incentives, and assess their consequences in terms of organizational economics.

Although peer production as a phenomenon is not totally new to economists, this dissertation is the first to introduce the concept and take Wikipedia and Open Source Software as particular instantiations of an emerging organizational model,¹ which ought to be sustainable whenever (i) the production process is highly modular (i.e. it can be divided into small and independent sub-tasks), (ii)

¹ Following-up on Lerner & Tirole’s (2002) seminal paper, which asks the question of “why would thousands of top-notch programmers contribute freely to the provision of a public good”, the overwhelming majority of the existing literature has focused its attention on rationalizing the success of Open Source Software and concludes that the model is unlikely to extend beyond software production. (See, e.g., Lerner & Tirole (2005). Maurer & Scotchmer (2006) is an exception). It is informative in this respect to note that only two economics papers have been published on Wikipedia since its inception in 2001, compared to thousands in the other social sciences.

the modules are intrinsically motivating and (iii) production does not require high capital inputs (as communities of volunteers seem good at collectively organizing labor but not managing capital).²

Because peer production, when sustainable, essentially eliminates contracting costs and agency problems (as contributors are intrinsically motivated to work), and accelerates discovery through automatic disclosure, it can benefit from a competitive advantage over proprietary production models whenever the production process requires input from a highly diverse and skilled labor force.³

Alternatively, it could also be successfully applied to any cumulative innovation process that needs to draw upon a wide variety of skills (see, e.g., Maurer & Scotchmer (2006) for the example of drug discovery and innovation in the biotech sector). One implication of the above argument is that, in a cumulative innovation framework, peer production will benefit from a significant comparative advantage for developing breakthrough innovations, as there is no hierarchical management structure to discourage directions that are highly uncertain and/or not immediately profitable.⁴

Roadmap of the dissertation

The first contribution of this dissertation is a methodological one. **Chapter 1** reports on the development of an online experimentation platform specifically designed to strengthen the internal validity of the decisions elicited over the Internet. The platform provides controls over many of the confounding factors that could prevent experimentalists from running experiments over the Internet. In particular it (i) controls for differences in response times, (ii) deals with the issues of selective attrition, concentration and distraction and (iii) provides as much control as possible over subjects' beliefs as regards the experimental instructions. The methodology is applied to the elicitation of social and risk preferences within a sample of traditional laboratory subjects. Since the decision interface is usable as is in the lab (through an Internet browser), subjects are randomly assigned to either a "laboratory" or an "Internet" condition (i.e. at home) for comparison purposes. All in all, using the same subject pool, the same decision interface and the same monetary stakes, the comparison concludes in favor of the reliability of behaviors elicited through the Internet according to the additional controls of the design.

Chapter 2 of this dissertation takes Wikipedia as one paradigmatic example of peer production, in which extrinsic motivations play no role in shaping contributors' behavior.⁵ It relies on this ideal study site to provide the first comprehensive field test of the relative role of each class of social motive that economic theory has put forward to account for people's willingness to sustain cooperation in public goods like environments. Indeed, competing models based on altruism, reciprocity and social

² I am currently in the process of formalizing the sustainability conditions and establishing the nature of peer production's competitive advantage over proprietary production models in an upcoming theoretical paper.

³ As a result, beyond software, peer production can also be successfully applied to any information good. A case in point is the online encyclopedia Wikipedia, which, since its inception in 2001, has both achieved high reliability (Giles 2005) and put Encarta and Britannica, its for-profit competitors, out of business.

⁴ This theoretical argument is very similar to the one developed by Aghion et al. (2005) and empirically supported by Williams (2013) according to which early-stage innovation is better conducted by academic researchers rather than corporate ones. Indeed, some early stage lines of research may never get started in the private sector, as those may not look viable according to a net present value criterion.

⁵ Here I define as extrinsic a motive that is monetary in nature, be it in the short (immediate payment) or in the longer run (labor market signaling). See chapter 2 of this dissertation for further details on how Wikipedia works in practice.

image preferences have been tested extensively in the laboratory, but there is very little field evidence as to which of those matter most in economically relevant contexts.

By doing so, this chapter illustrates the research benefits that accrue when coupling experimental and computational methods in an online context. While it is possible to reliably elicit subjects' altruism and reciprocity preferences with an online experiment, it has generally been a challenge to elicit social image motives experimentally (even more so in a decontextualized fashion). To achieve this goal, the chapter relies on the wealth of observational data available from Wikipedia on contributors' behavior to construct individual measures of revealed preference for social image within the Wikipedia community. Those measures rely on (i) contributors' propensity to post more or less information about themselves on their Wikipedia user page (which is of no direct use to efficiently contribute to the encyclopedia) and (ii) their propensity to prominently display the signs of social recognition that they received from other contributors to the entire community. It is then possible to relate subjects' preferences to detailed real-world records of their contributions to the Wikipedia project, which are separately extractable from the website.

Based on a representative sample of 850 Wikipedia contributors, the chapter reports that reciprocity and social image are both strong motives for sustaining cooperation in this public goods like environment, while altruism seems to play less of a role. This result strikingly confirms the conclusions of the existing laboratory literature on the private provision of public goods. An important result is that across all specifications, reciprocity and social image consistently appear as substitutable motivational drivers rather than complementary ones. The chapter also reports on the specific patterns of contributions of the Wikipedia administrators, a specific group of contributors who self-selected into performing a policing role within the community and are notably in charge of dealing with disruptive users. While trust in anonymous strangers (as measured by a standard Trust game) is unrelated to contribution levels among regular contributors, the estimates show that administrators who are less trusting are significantly more active and more likely to exercise their policing rights.

Chapter 3 of the dissertation makes a contribution to labor economics by bringing experimental economics within real-world productive organizations. The chapter focuses on the community of Open Source Software developers (OSS) and elicits their social motivations with an online experiment. Similar to the Wikipedia case, the online activities and contributions of individual developers are tractable at a very detailed level, which is rarely the case in traditional corporate environments.

The combination of experimental and field data on the community of OSS developers provides an opportunity to address a number of longstanding but reputedly difficult questions in the literature. It notably allows to study how heterogeneous motivations affect the extent and nature of individual contributions to team level efforts. Importantly, the combination of experimental and field data allows to test for endogenous sorting behavior by social type at the team level. Indeed, while there is substantial theoretical and experimental evidence that more cooperative types seek to match assortatively within groups and organizations in order to sustain high cooperation levels (which typically has a significant impact on efficiency), the difficulty to collect the necessary data has prevented the existing literature from testing for such endogenous sorting behaviors in the field. Last,

the approach also allows to assess the impact of team composition on the likelihood of success at the project level.⁶

Two features of the OSS community are essential to perform the above tests and distinguish it from the community of Wikipedia contributors: (i) many self-formed development teams working on separate software projects simultaneously coexist and (ii) approximately half of the developers who contribute to OSS actually derive some monetary payment from their contributions.⁷

Based on a stratified sample of 1,194 OSS developers, the chapter reports that social motivations predict developers' contributions as strongly as extrinsic ones, but have a different impact on the nature of participation. Socially motivated developers, tend to join less development teams (i.e. they have a lower extensive margin of participation) but contribute significantly more to each (with an overall positive association with contributions). The chapter also reports strong evidence of endogenous team-level sorting by cooperative type within the sample of OSS developers. Free-riders seek to join development teams that are comprised of different social types than their own, while strongly cooperative developers tend to match assortatively. This result provides the first field validation of a consistent laboratory finding in the experimental literature on the private provision of public goods. The assortative matching effect seems to be largely driven by reciprocating and altruist project administrators, who ultimately get to choose who joins to development team and therefore act as the gatekeepers of their teams by seeking to coopt developers of their own cooperative type.

Figure 1 summarizes the overall logic of the present work, as well as the contribution of this dissertation to the literature. Chapter 1 develops and methodologically assesses the reliability of a novel tool for running online economic experiments. Chapter 2 and 3 document the rise of a new and significant model for organizing production – peer production – in which agents voluntarily self-assign work and successfully coordinate towards the provision of global public goods without necessarily relying on monetary incentives. Those chapters then use the context of peer production and rely on a combination of experimental and computational methods to respectively (i) provide the first comprehensive field test of the theory of the private provision of public goods, (ii) study the importance of social preferences as work motives within real-world productive organizations and (iii) report the first field evidence of endogenous sorting behavior of economic agents within productive teams based on their cooperative types.

⁶ Chapter 3 does not report on this latter topic however, as the necessary data is currently being collected.

⁷ Many firms support the development of OSS projects directly, notably by allocating a fraction of their labor force to their development. See chapter 3 of the dissertation for further details.

References

- Aghion, P., Dewatripont, M. & Stein, J., 2005. *Academia, the private sector, and the process of innovation*, National Bureau of Economic Research.
- 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.
- Ghosh, R. A., 2007. *Economic impact of open source software on innovation and the competitiveness of the Information and Communication Technologies (ICT) sector in the EU*. Maastricht: UNU-MERIT.
- Giles, J., 2005. Internet encyclopaedias go head to head. *Nature*, 438(7070), pp.900–901.
- Greenstein, S., & Nagle, F., 2014. Digital Dark Matter and the Economic Contribution of Apache. *Forthcoming, Research Policy*.
- Henrich, J., Heine, S.J. & Norenzayan, A., 2010. The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), pp.61–83.
- Lazer, D. et al., 2009. Life in the network: the coming age of computational social science. *Science*, 323(5915), pp.721–723.
- Lerner, J. & Tirole, J., 2002. Some Simple Economics of Open Source. *The Journal of Industrial Economics*, 50(2), pp.197–234.
- Lerner, J. & Tirole, J., 2005. The Economics of Technology Sharing: Open Source and Beyond. *The Journal of Economic Perspectives*, 19(2), pp.99–120.
- Levitt, S.D. & List, J.A., 2009. Field experiments in economics: The past, the present, and the future. *European Economic Review*, 53(1), pp.1–18.
- Loewenstein, G., 1999. Experimental Economics From the Vantage-point of Behavioural Economics. *The Economic Journal*, 109(453), pp.25–34.
- Maurer, S.M. & Scotchmer, S., 2006. *Open Source Software: The New Intellectual Property Paradigm*, National Bureau of Economic Research.
- Walli, S., Gynn, D., & Rotz, B., 2005. *The growth of open source software in organizations*. Report.
- Williams, H.L., 2013. Intellectual Property Rights and Innovation: Evidence from the Human Genome. *Journal of Political Economy*, 121(1), pp.1–27.

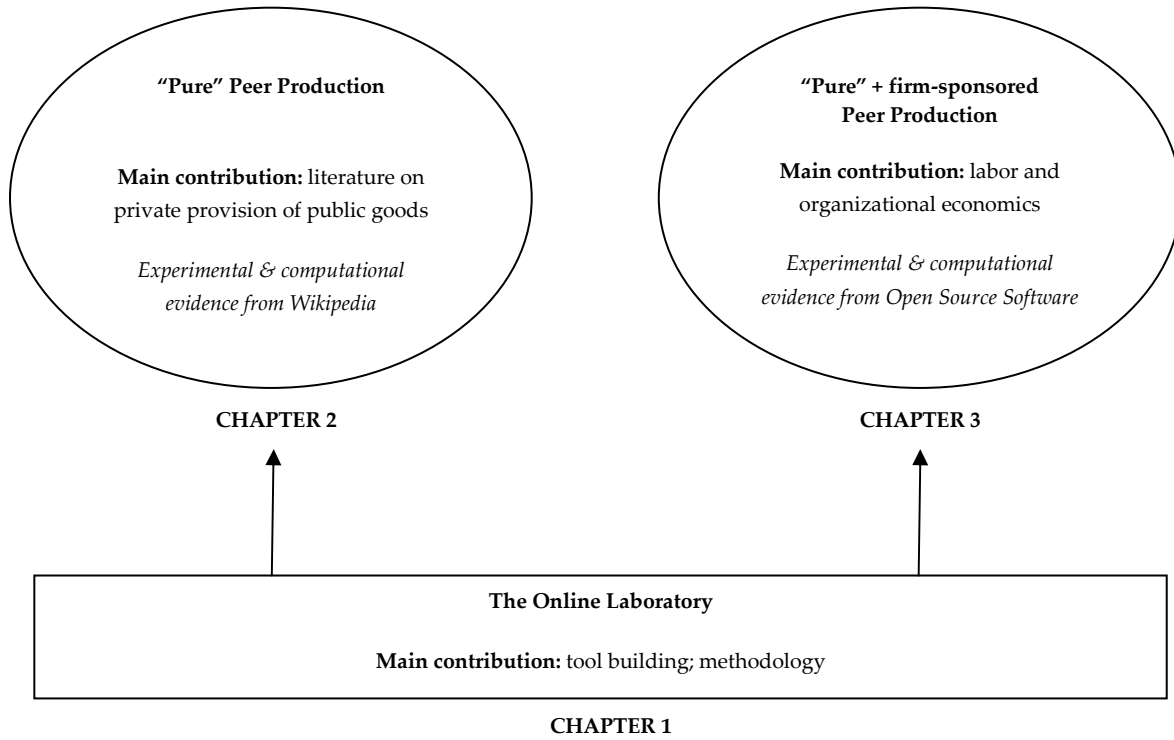
Tables and figures

Table 1. World Internet users, growth and penetration statistics

World region	Total Population (June 30, 2012)	nb of Internet users (Dec. 31, 2000)	nb of Internet users (June 30, 2012)	Internet penetration rate in % (June 30, 2012)	nb of Internet users growth rate in % (2000-2012)
Africa	1,073,380,925	4,514,400	167,335,676	15.6	3,606.7
Asia	3,922,066,987	114,304,000	1,076,681,059	27.5	841.9
Europe	820,918,446	105,096,093	518,512,109	63.2	393.4
Middle East	223,608,203	3,284,800	90,000,455	40.2	2,639.9
North America	348,280,154	108,096,800	273,785,413	78.6	153.3
Latin America	593,688,638	18,068,919	254,915,745	42.9	1,310.8
Oceania	35,903,569	7,620,480	24,287,919	67.6	218.7
World total	7,017,846,922	360,985,492	2,405,518,376	34.3	566.4

Source: www.internetworldstats.com

Figure 1. The logic of the dissertation



CHAPTER 1

**SOCIAL PREFERENCES IN THE ONLINE LABORATORY:
A RANDOMIZED EXPERIMENT**

Social Preferences in the Online Laboratory: A Randomized Experiment*

Jérôme Hergueux[†]

Nicolas Jacquemet[‡]

May 15, 2014

Forthcoming: Experimental Economics

Abstract

Internet is a very attractive technology for the implementation of experiments, both in order to obtain larger and more diverse samples and as a field of economic research in its own right. This paper reports on an experiment performed both online and in the laboratory, designed to strengthen the internal validity of decisions elicited over the Internet. We use the same subject pool, the same monetary stakes and the same decision interface, and control the assignment of subjects between the Internet and a traditional university laboratory. We apply the comparison to the elicitation of social preferences in a public good game, a dictator game, an ultimatum bargaining game and a trust game, coupled with an elicitation of risk aversion. This comparison concludes in favor of the reliability of behaviors elicited through the Internet. We moreover find a strong overall parallelism in the preferences elicited in the two settings. The paper also reports some quantitative differences in the point estimates, which always go in the direction of more other-regarding decisions from online subjects. This observation challenges either the predictions of social distance theory or the generally assumed increased social distance in internet interactions.

JEL classification: C90, C93, C70

Keywords: Social Experiment, Field Experiment, Internet, Methodology, Randomized Assignment

* This paper is a revised and augmented version of CES Working Paper n° 2012-70. We are grateful to Anne l'Hôte, Andrews-Junior Kimbembe and Ivan Ouss for their outstanding research assistance, as well as Maxim Frolov and Joyce Sultan for their help in running the laboratory experimental sessions. We are especially indebted to Yann Algan for his help during the development of this project. We gratefully thank the editor, Jacob Goeree, two anonymous referees and Guillaume Fréchette, Olivier L'haridon, Stéphane Luchini, David Margolis, Ken Boum My, Paul Pézanis-Christou, Dave Rand, Al Roth, Antoine Terracol, Laurent Weill and the members of the Berkman Cooperation group for helpful remarks and discussions. We also thank seminar participants at the Berkman Center for Internet & Society at Harvard and the 2012 North American Economic Science Association conference for their comments. We gratefully acknowledge financial support from the European Research Council (ERC Starting Grant). Jacquemet acknowledges the *Institut Universitaire de France*.

[†] University of Strasbourg, Institute of Political Studies and Sciences Po, Department of Economics. Research fellow, Berkman Center for Internet & Society at Harvard University. 23 Everett Street, 2nd floor, Cambridge, MA, 02138, USA. e-mail: jhergueux@cyber.law.harvard.edu

[‡] Université de Lorraine (BETA) and Paris School of Economics. 3 Place Carnot, 54035 Nancy. e-mail: Nicolas.Jacquemet@univ-lorraine.fr

1 Introduction

In the field of experimental economics, it is a long time since researchers called for the development of the “online laboratory” (Bainbridge 2007). The interest in online experimentation has been propelled by the possibility of reaching more diverse samples, recruiting larger subject pools and conducting cross-cultural social experiments in real time at an affordable cost.¹ Besides this methodological concern, the Internet is becoming an increasingly prominent experimental field for social science research in its own right (see, e.g., Resnick et al. 2006; Chesney et al. 2009), as we live more and more of our social and economic lives online. It is thus essential to conduct experiments directly over the Internet if we are to rely on the experimental method to understand the various types of social and economic activities that people engage in online.

Notwithstanding these appealing features, the development of the “online laboratory” still remains in its infancy. The primary goal of this paper is to help fill this gap by conducting a methodological evaluation of an Internet-based experimentation procedure. Horton et al. (2011) underline the difficulty of coming up with procedures for online experiments that ensure their internal validity, *i.e.* the possibility of confidently drawing causal inferences from one’s experimental design. A number of confounding factors have been identified that have probably prevented researchers from running experiments online: (i) it is difficult to monitor the identity of subjects participating in the experiment, (ii) subjects may read the experimental instructions too carelessly and/or make decisions too quickly and/or get significantly distracted during the course of the experiment, (iii) subjects may selectively drop out of the experiment in ways that the experimenter does not understand, (iv) subjects may not believe that they interact with other human players and/or that they are going to be paid at the end of the experiment as described in the instructions, and finally (v) the issue of reliably and automatically processing the payment of subjects over the Internet in an anonymous fashion appeared to be a major blocker.

In this paper, we seek to compare the behavioral results generated both in a traditional laboratory and over the Internet. To do so, we develop an online platform specifically dedicated to conducting social experiments over the Internet that is usable as in the laboratory. To account for the effect of self-selection between implementations, we control the allocation of subjects between treatments.

The platform provides controls over many of the above-mentioned confounding factors. In particular we (i) control for differences in response times, (ii) deal with the issues of selective attrition, concentration and distraction and (iii) provide as much control as possible over subjects’ beliefs as regards the experimental instructions.

The existing literature has already covered a variety of different games implemented over the Internet (Table 1 summarizes the methodology and main conclusions of this literature). The seminal study of Anderhub et al. (2001) focuses on an individual level decision experiment under uncertainty, both in the laboratory and online. Shavit et al. (2001) compare student bids over buying prices for simple lotteries both in the classroom and online. Charness et al. (2007) also compare classroom

¹ In a recent paper, Henrich et al. (2010) warned against behavioral scientists’ current over-reliance on data overwhelmingly gathered from populations of Western undergraduate students and recommended a major effort to broaden the sample base. The Internet is a promising medium for conducting experiments with large and diverse samples. It is now possible to reach 78.3% of the North American population through the Internet, and while only 11.4% of the African population can currently be reached through this method, the exponential growth of its user base (from 4 million users in 2000 to 118 million users in 2011) could soon make it an attractive tool for conducting experiments in the developing world as well (*source*:www.Internetworldstats.com).

experiments with other Internet-based experimental settings to investigate the effect of social distance on trust and reciprocity in a simple lost wallet game. They find that trust and reciprocity both decrease in an Internet-based setting, which they argue is consistent with social distance theory (Akerlof 1997). Fiedler and Haruvy (2009) and Chesney et al. (2009) take an exploratory approach and build a virtual laboratory on the *Second Life* website. Chesney et al. (2009) recruit subjects from the *Second Life* community to perform a series of social experiments and compare the results with those of the traditional laboratory literature. Similarly, Fiedler and Haruvy (2009) recruit subjects from *Second Life* to perform a Trust game, but directly compare their results with those obtained from traditional laboratory subjects playing in the same virtual environment, but in a physical laboratory. They also find trust and trustworthiness to be lower outside the physical lab. Most recently, Horton et al. (2011) and Amir et al. (2012) have used the online labor market platform *Amazon Mechanical Turk* to conduct a set of classic experiments and replicate qualitatively some general results drawn from the experimental economics literature.

We contribute to this burgeoning literature by looking at social preferences and by providing a rigorous comparison of the Internet-based experimentation with traditional lab experiments. We apply our methodology to the measurement of social preferences – combined with a risk aversion task – through a Public Good game, a Trust game, a Dictator game and an Ultimatum game (using a within-subjects design). The main conclusions that we draw from this comparison are twofold. First, the social preferences elicited in the lab and online are qualitatively very similar – all common inferences on social preferences that we replicate in the laboratory would also be obtained based on online data. Second, we do, however, observe some differences in the point estimates between treatments. Social distance theory (Akerlof 1997) predicts that the stronger anonymity that prevails in Internet-based interactions should drive social preferences down as compared with the laboratory setting, where people can (i) see each other before and after the experiment, (ii) recognize that they often come from the same socio-economic background and (iii) know that they are going to be matched with one another during the experiment. On the contrary, we find robust and significant evidence that subjects allocated to the Internet treatment behave more altruistically and, when insignificant, the differences in social preferences always go in the direction of more other-regarding decisions online. We suggest an explanation for our results based on the nature of the social and economic interactions in which individuals tend to engage online, which they are likely to bring to the experiment through its contextual implementation.

Our results are important to the community of researchers wishing to develop the online laboratory as a medium for running social experiments over the Internet and to relate their results to the established laboratory literature. They are also important for social scientists wishing to use social experiments to research the Internet as a field: given the observed parallelism between fields, it makes sense for researchers to bring their experimental tools directly to the field, *i.e.* over the Internet, if they want to learn from subjects' behavior in this context, rather than sticking with the more difficult approach of trying to bring a subsample of those subjects into a traditional university laboratory.

The rest of the paper proceeds as follows. Section 2 documents the design of the experiment, reports on the development of our online experimental economics platform and explains our experimental procedures. Section 3 reports the main results of the experiment. Section 4 provides additional evidence on the reliability of the comparison based on an analysis of the internal validity of the online experiment, secondary outcomes and robustness treatments. We discuss the main outcomes of this comparison in Section 4, and conclude in Section 5.

2 Design of the experiment

Social isolation and greater anonymity are well-recognized distinctive features of online interactions. In order to provide a rather conservative testbed comparison between online and lab experiments, we focus on the elicitation of social preferences. Shavit et al. (2001) have also shown that subjects tend to be less risk-averse when making decisions online rather than in a classroom. We thus complement our preference measures with a risk aversion task. Our main methodological contribution is to build an Internet-based experimental environment which can be implemented both online and in the laboratory. We conclude this section with a detailed description of the procedures and decision interface we used.

2.1 The decision problems

At the beginning of the experiment, each subject is attributed a role: either participant A or participant B. The assigned role remains the same during the whole experiment. The experiment is divided into two different parts. First, we elicit decisions in five different games. The first four games are taken from the social preferences literature (see, e.g., Fehr & Camerer 2004) while the last one elicits individual risk aversion. At the end of each game, subjects are asked to answer non-incentivized questions about their beliefs and intentions in the game they have just played. In the second part of the experiment, subjects are asked to answer some standard demographic and social preference-related questions, along with some questions eliciting their beliefs about the study.

Public Good Game. 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. Following Fischbacher et al. (2001), we elicit both unconditional and conditional contributions, asking subjects to make two contribution decisions in turn. They first decide on how much of their 10€ they want to invest in the common project. They then 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.² One of the two decisions is randomly drawn to be binding and determines the individual earnings for this game according to the following payoff function:

$$\pi_i = 10 - \text{contrib}_i + 0,4 \sum_{j=1}^4 \text{contrib}_j \quad (1)$$

Right after the decision screen, we ask subjects about (i) their normative opinions about how much people *should* contribute to the public good (ii) whether they had an idea about how much the other

² The second decision is a variant of the “strategy method” (Selten 1967), introduced by Fischbacher et al. (2001) to elicit conditional cooperation. As in the original strategy method, subjects are asked decisions for each possible state of the world, but these states are reduced to average contributions of other subjects instead of all possible combinations of their individual decisions. In order to give subjects a monetary incentive to take both decisions seriously, we applied the same compensation rule as in Fischbacher et al. (2001): for one randomly chosen subject, the table of unconditional decisions is binding; for the other three the relevant decisions are the unconditional ones. These realizations of the draw are the monetary outcomes of this stage for each subject.

members of their group would contribute to the public good when they made their decision, and if so (iii) their beliefs about how much the other members of their group actually contributed on average.

Dictator Game. Each participant A is matched with a participant B and plays the role of dictator. The dictator receives a 10€ endowment, of which he must decide how much is transferred to participant B. The difference is participant A's earning for this game.

Ultimatum Bargaining Game. Each participant A is matched with a participant B. Participant A is the proposer and must decide on how much of an initial endowment of 10€ is transferred to participant B – the responder. The responder is simultaneously asked for the threshold level of transfer below which the offer will be refused. The earnings of each player in this game are computed according to the proposal if participant A's transfer is higher or equal to the threshold. Otherwise, both players' earnings are set equal to 0.

Trust Game. 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 the trustor. 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. Right after the decision screen, we ask trustors about (i) whether they had an idea about how much the trustee would return to them when they made their decisions, and if so (ii) their beliefs about the amount that the trustee would return.

Risk aversion elicitation. Each participant faces a menu of ten choices between lottery pairs, adapted from Holt & Laury (2002). The probability of getting the higher amount is always the same between the two lottery pairs, but the safe option pays either 20€ or 16€ while the risky option pays either 38.5€ or 1€. The probability that subjects get the higher amount in both options steadily increases from 10% in the first decision problem to 100% in the last one. Thus, in decision 10, subjects actually choose to earn either 20€ or 38.5€ with certainty. One of the ten decisions is randomly drawn to determine the binding lottery choice. Earnings for this game are then derived from a random draw according to the probability of the corresponding lottery.

Social values survey. After all games have been played, subjects are asked to fill in a questionnaire with some standard demographic questions followed by social preference-related questions. This set of questions has been taken from the *World Values Survey* (WVS), the *General Social Survey* (GSS) and the *German Socio-Economic Panel* (GSEP) – the three most commonly used sources in the empirical literature. Specifically, we ask subjects:

(i) to what extent they consider it justifiable to free-ride on state benefits (cooperation variable; WVS question);

(ii) whether they think that people are mostly looking out for themselves as opposed to trying to help each other (altruism variable; WVS question);

(iii) whether they think that people would try to take advantage of them if they got a chance as opposed to trying to be fair (fairness variable; WVS question);

(vi) whether they think that most people can be trusted or that one needs to be very careful when dealing with people (trust variable; WVS and GSS question);

(v) how trusting they generally are of people (trust variable; GSEP question);
(vi) how trusting they are of people they have just met (trust variable; GSEP question);
(vii) whether they generally see themselves as fully prepared to take risks or as trying to avoid them (a question taken from Dohmen et al. 2011). All questions are mandatory and none is remunerated.

Debriefing questionnaire. As demonstrated by Eckel and Wilson (2006), the internal validity of online experiments can be challenged by subjects' skepticism about whether they actually interact with other human subjects and whether they will actually be paid according to the rules described in the instructions. To get some control over these dimensions, we ask subjects to rate their level of confidence in those two critical features of the study. As a complement, we end the survey by asking subjects to report on how carefully they read the experimental instructions, on how calm their environment was when they performed the experiment and on whether they had participated in any similar studies in the past.

2.2 Procedures common to both implementations

All five games, followed by the survey, are played successively in each experimental session. As we seek to elicit social preferences in isolation from learning effects and strategic concerns, each game is only played once. To neutralize reputation effects, we match subjects in each game according to a perfect stranger procedure. Last, in order to further break any possible correlation between games, only one game out of the whole session is randomly drawn as binding to compute each subject's earnings. Final payoffs equal the earnings from the corresponding decision plus a 5€ show-up fee. Subjects are only informed of their earnings in each game at the very end of the experiment.

As all games are played one after the other, order effects could influence the preferences we elicit. This led us to implement three different orderings. The Public Good game is the most cognitively demanding, so we start all sessions with this game. The Dictator, Ultimatum and Trust games all appear afterwards in varying orders. As we mainly use the risk aversion task for purposes of replication and as a control variable, we maintain this decision problem as the last in all sequences.

- Order 1: *Public Good – Dictator – Ultimatum – Trust – Risk Aversion*
- Order 2: *Public Good – Trust – Ultimatum – Dictator – Risk Aversion*
- Order 3: *Public Good – Ultimatum – Dictator – Trust – Risk Aversion*

Subjects face the exact same decision interface both in the lab and online. 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.³ The first screen of the decision interface provides subjects with general information about the experiment, including the number of sections and how their earnings will be computed. Each game is then performed in turn, following a given sequence of screens.

³ The interface has been developed under *Lime Survey* (<http://www.limesurvey.org/>), a highly customizable open-source survey tool.

The first screen of each section describes the instructions for the game that subjects are about to play (Figure 1 provides an English translation of the original instructions in French for the Trust game).

One important methodological concern with online experiments is to guarantee an appropriate understanding of the decision problems when no interaction with the experimenter is possible, which makes it difficult, for instance, to rely on the standard post-instructions questionnaire coupled with oral questions. We address this issue through several distinctive features of the interface. First, we include suggestive flash animations illustrating the written experimental instructions at the bottom of each instruction screen (the animation appears at the bottom of the first screen, as shown in Figure 1; the animation is illustrated in Figure 2 by step-by-step screen captures).

Displaying a purely random sequence of flash animations would introduce uncontrolled and subject specific noise – through, e.g., anchoring on a particular behavior or sequence of events.

Since our main objective is to compare behavior between the two implementations, we get rid of this noise by fixing the actual sequence: the loop of concrete examples displayed in the animations is 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.

Second, the instruction screens are followed by a screen providing some examples of decisions, along with a 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 Good and Trust games (English translations of the original earnings calculators in French are provided in Figure 3, (a) for the Public Good game and (b) for the trust game). In contrast to the flash animations, the numeric results of each scenario run by a subject in the earnings calculator screens are explicitly displayed.

Last, the system provides 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 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 (Figure 4 provides an English translation of the original decision screen in French for the Public Good game).

A potentially important confound when comparing laboratory and online experiments is the average variation in decision times. Anderhub et al. (2001) report that subjects make decisions more quickly in an online environment. However, an established body of research in psychology indicates that shorter decision times are likely to be associated with *instinctive and emotional reasoning processes* rather than *cognitive and rational* ones (Kahneman 2003), which could cause subjects to make more pro-social decisions on average. In order to generate a control variable for this dimension, the platform recorded detailed data on the time in seconds that subjects spent on each screen of the interface (this timer was not visible to the subjects). But more time on a screen does not necessarily mean longer decision time if, for instance, online subjects leave their computer while answering the survey.

To get further information about whether some subjects were likely to have been distracted from the online experiment at some point, we included an indicator of mouse inactivity in the platform.

The indicator records both the screen and the duration of inactivity each time the mouse of the subject is inactive for more than 5 minutes.⁴

2.3 Practical implementation of the experiment

All participants in the experiment were contacted through the subjects database of the experimental economics laboratory of University Paris 1 Panthéon-Sorbonne.⁵ The allocation to sessions is intended to minimize differences in the subject pools and avoid self-selection into treatments. We apply a matching procedure that proceeds in two steps. First, subjects are invited to register for a date on which a session takes place. They are told that practical details about the experiment will follow once their registration has been confirmed (as usual, registrations are confirmed on a first-come first-served basis). Indeed, two sessions are scheduled during each time slot: one session online and one session in the laboratory. In the second step, we sequentially allocate subjects either to the laboratory or to the online experimentation according to their registration order.

As the capacity of the laboratory allows for no more than 20 subjects, we allowed 56 persons to register for each time slot, allocating half of them to the laboratory and the other half to the Internet session. In the laboratory, we had to refuse any overbooked subjects who showed up on time. Since no such constraint applied to the online experiment, we allowed all subjects to participate while keeping track of those who logged-in after the target number of 20 participants had been reached. In laboratory sessions, subjects are randomly assigned to a computer upon arrival. The instructions for the experiment are read aloud, and subjects are then left to use all devices at their disposal to check their own understanding (access to the text, earnings calculators, etc.). Each game is described in turn, following the above-described interface, so that all subjects progress inside the experiment at the same time.

Online subjects are invited to visit the url embedded in their confirmation e-mail at the time their session is scheduled, and to log into the system using their e-mail address, which served as a unique login token. The url was activated during the half-day spanning the time scheduled for the experiment. The computer program allocates online subjects to either participant A or participant B according to their login order (in order to ensure that we get a somewhat equal split of the subject pool between participant As and participant Bs, despite possible dropouts).

At the end of the experiment, subjects are matched using a perfect stranger procedure. Subjects are informed of their earnings in each game only at the end of the experiment. In the laboratory, subjects from a given session are matched together. By contrast, online subjects had their decisions matched with the decision records of subjects who had already completed the experiment.⁶ This feature of the platform allowed Internet subjects to perform the experiment independently and at their own pace, thus smoothing the interactions and arguably reducing dropouts.⁷ The drawback of this matching procedure is that it breaks the joint determination of payoffs between subjects: when a subject makes a decision, his own payoff is determined by the decision made by some previous participant, while

⁴ The system considered the mouse inactive when it was moving over screens not belonging to the experimental economics platform.

⁵ The database is managed using Orsee (Greiner 2004).

⁶ Since we apply a sequential matching rule for online subjects, the queue has to be initialized somewhere. We used data from 3 pilot sessions in the laboratory run during summer 2010 in preparation for the current study.

⁷ Overall, 208 subjects logged in to the platform to participate in the online experiment, of whom 6 dropped out before completion.

his current decision determines the payoff of another, future participant. Such a sequential matching between current and past decisions can hardly be avoided in online experiments, in which subjects must be allowed to participate at any time they see fit. An alternative way of implementing the online matching, introduced by Cooper and Saral (2013), would have been to compute both subjects' outcomes at a later time, once the second subject has gone through the experiment – thus restoring the joint determination of payoffs inside each pair. We opted for the first solution for two reasons. First, having subjects wait until a future date before they can get their earnings involves inter-temporal preferences and may induce further differences in the saliency of payoffs between the two environments. Second, we were also concerned that the credibility of the experiment would be challenged for online subjects, if they were not informed about their experimental earnings immediately after their participation. Both solutions have advantages and drawbacks, and a more systematic comparison of the consequences of each design is worth investigating in future works.⁸

Laboratory subjects' earnings are paid in cash before subjects leave the laboratory. Internet subjects get paid through an automated PayPal transfer. This guarantees a fungibility similar to that of cash transfers, as money transferred via PayPal can be readily used for online purchases or easily transferred to one's personal bank account at no cost. 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.

2.4 Summary of the design

To sum up, the experiment elicits the same decisions with similar procedures in both treatments. In particular, we recruit from the same subject pool, use the same monetary stakes, the same decision interface, and control the allocation of subjects between the lab and Internet treatments. This is summarized in Table 2.1, which also provides an exhaustive list of all the preferences we elicit.

At the same time, there are some important practical differences between the two kinds of implementations, most of which are due to subjects not being in the same physical space as the experimenter in the online implementation. Obviously, the standard procedure for laboratory experiments does not have to be adapted to such constraints. Our empirical strategy is to stick to common practice with the laboratory implementation, so as to keep the benchmark situation as close as possible to existing evidence. We tried to choose the most innocuous adaptations when we had no choice but to introduce a difference between the two designs. Table 2.2 summarizes the resulting differences between our two treatments.

We conducted two different sets of experimental sessions, each conducted over a one-week period: 6 sessions (3 in the lab, 3 online) were conducted in November 2010 and 12 sessions (6 in the lab, 6 online) were conducted in November 2011.⁹ Overall, 180 subjects performed the experiment in the laboratory and 202 subjects performed it online. We conducted 8 sessions with games order 1 (80 participants in the lab, 85 online), 6 sessions with games order 2 (60 and 67) and 4 sessions with games order 3 (40 and 50). Subjects in both conditions earned on average 21.24€ from the experiment.

⁸ Our robustness treatments, presented in Section 4.3, provide some preliminary insights on this issue.

⁹ The 2010 version of the experimental economics platform did not elicit subjects' level of confidence in the experimental instructions, nor did it collect detailed data on the time spent by subjects on each screen of the interface. After observing that overall response times did indeed significantly differ between treatments, we decided to include those features before conducting further sessions.

3 Social preferences in the online laboratory

This section reports on our main outcome of interest, *i.e.* the reliability of the online elicitation of social preferences, taking laboratory behavior as a benchmark. In the next section, we assess the internal validity of both the online experiment and the comparison with laboratory behavior, based on the analysis of underlying secondary outcomes and additional robustness treatments.

Figure 5 provides a qualitative comparison of the behavioral patterns observed in the lab and online. For all games, the preferences we elicit online are parallel to those generally observed in the laboratory – which our lab condition replicates. While the theoretical prediction in the Public Good game is full free-riding, we do observe a positive amount of contribution that ranges between 35% and 40% of the initial endowment. In particular, the Nash equilibrium of the one-shot game is strongly rejected everywhere, with a high share of subjects making other-regarding decisions.

In the Dictator game (Figure 5.g), we observe three striking variations when preferences are elicited online. In the laboratory, the mode of the distribution is at 0, with 40% of subjects deciding not to give anything to their partner. For behavior online, the share of zero donors falls to half of this proportion and the mode of the distribution is equal to 5 (*i.e.* equal split). Last, at the upper tail of the distribution, some subjects are willing to send more than 70% of their endowment online while no such behavior is observed in the laboratory. All three inflexions go in the direction of more other-regarding decisions online. In the Ultimatum Bargaining game (Figure 5.e), the shape of preferences for proposers are much more parallel, although we still observe a slightly higher share of zero donors in the laboratory (5%) as compared to online subjects (0%). Similarly, for receivers (Figure 5.f), the observed patterns are very similar with a mode at the equal split threshold, although there exists a slight difference at the bottom of the distribution with the share of low thresholds being 5% higher in the laboratory.

In both the Trust game (Figure 5.c) and the Public Good game (Figure 5.a), the same qualitative variation as in the Dictator game can again be observed: the high share of non-participants in the laboratory (1/4 of senders in the trust game, 1/5 in the Public Good game) is strongly reduced online, falling to around 1/10 in both instances. The remaining shape of the distribution is comparable, which again tends to suggest that players tends to be more pro-social online. Figures 5.b and 5.d describe the decisions elicited through the strategic method. Figure 5.b focuses on the Public Good game and plots the mean of the contributions to the common project made by subjects in the laboratory and Internet conditions, conditional on the average contribution made by the other 3 group members. In both fields, the qualitative pattern is very similar, with conditional contributions that are monotonically increasing in the average contributions of others but with a slope that is strictly lower than one. As this average group contribution increases, the distribution of conditional contributions among Internet subjects tends to dominate the distribution of conditional contributions among laboratory subjects, potentially indicating that online subjects were more prone to conditional cooperation. The overall effect, however, is relatively weak.

Figure 5.d, by contrast, exhibits a much stronger pattern. It plots the mean of the amount returned by participants Bs under laboratory and Internet conditions depending on the amount transferred by participant A. The shape of the social preferences elicited both online and in the laboratory points to the same conclusion: the amount returned by the trustee is strictly increasing in the amount received.

The slopes, however, are quite different. The distribution of returns among Internet subjects strictly dominates the distribution of returns among laboratory subjects.

One consistent result in the literature about Trust games is that trustors are generally willing to place some of their resources in the hands of trustees. For their part, trustees typically tend to exhibit positive reciprocity, but the effect is usually not strong enough for this to be profitable to the trustor (Fehr & Camerer 2004). We can see this general pattern in our data, whereby participants Bs exhibit positive reciprocity, but tend to systematically return a lower amount to participant As than they transferred in the first place. This result no longer holds among Internet subjects, however, in which participants Bs consistently return slightly more on average than the participant As initially transferred.

Last, regarding the risk aversion task, we follow Holt and Laury (2002) and interpret the number of times subjects chose the secure option as a raw measure of their level of risk aversion (Figure 5.h).¹⁰ Again, the overall patterns of risk aversion in each pool of subjects share the same qualitative features: very few subjects are observed at the lower end of the distribution. Most of the sample switches after 5 risky decisions, with the majority of subjects switching between decisions 5 and 9. The figure also shows, however, that the distribution of risk preferences online strictly dominates the distribution in the laboratory, indicating that levels of risk aversion tend to be lower online. This observation confirms the results reported in Shavit et al. (2001).

We now turn to a statistical assessment of the comparison. Table 3 reports on univariate non-parametric tests of differences between the two fields in terms of the mean and the dispersion of observed behavior. As regards mean comparisons, most of the differences discussed above induce statistically significant differences between the two elicitation fields (in 11 out of 14 measures). Leaving risk aversion aside, the most economically and statistically significant differences emerge in the Dictator game and the Trust game, especially as regards the behavior of trustees. On average, 58% of participant As in the Dictator game chose to transfer some fraction of their endowment to participant Bs in the lab, as opposed to 81% online. Overall, online subjects in the Dictator game transferred 17% more of their endowment to participant Bs. In the Trust game, they transferred about 9% more of their endowment, with this increase in trust being reciprocated in kind by participant Bs, who exhibited a reaction function to their transfers about 0.44 point steeper than laboratory subjects. Last, online subjects also appear significantly less risk-averse than laboratory subjects. The difference is significant at the 1% level, irrespective of whether we exclude confused subjects from the sample or not.

In their early experiments, Anderhub et al. (2001) and Shavit et al. (2001) both suggest that the variance in preferences tends to be higher when elicited online. Our statistical assessment does not confirm this conclusion. While the behavior in the Dictator game and risk aversion task do seem to be significantly more dispersed online, we actually find it to be significantly less dispersed for one of our measures of conditional cooperation in the Public Good game, and statistically indistinguishable from the variance generated in the lab for all the other measures.

Last, our risk aversion elicitation task allows us to directly investigate the issue of the quality of the data collected online. Overall, there were 13 inconsistent subjects in the laboratory as opposed to 44 online (two-tailed t-test, $p < 0.01$).

¹⁰ Note that in constructing this figure, we excluded from the analysis the 5 laboratory and 22 Internet subjects who arguably misunderstood the task and choose option A in decision 10. Apart from the last data point, including those subjects has no impact on the figure.

There was also a fair proportion of subjects who clearly misunderstood the task and chose option A in the last decision. 5 subjects did so in the laboratory, as opposed to 22 over the Internet (two-tailed t-test, $p < 0.01$). Consistent with previous findings, those results indicate that it is somewhat more difficult to obtain good quality data with web-based experiments, which should be compensated for by the ease with which the Internet allows to recruit larger samples.

To sum up, the comparison concludes that there is strong parallelism between the patterns of preferences elicited online and those elicited in a physical laboratory. We do observe some point differences between the two settings, though. Beyond the difference in risk attitudes (online subjects being less risk-averse), the most important differences in terms of social preferences are the intensity of altruistic behavior in the Dictator game and of the reciprocity of trustees in the Trust game. What is more, whether the differences are statistically significant or not, they always go in the direction of stronger other-regarding preferences when the elicitation takes place online. We now turn to additional evidence intended to assess the robustness of this surprising result as regards existing theories of social preferences applied to online environments.

4 Do subjects actually behave more pro-socially online?

To assess the robustness of our comparison, we first focus on factors that may impede the internal validity of our observations: composition effects in the subjects' pool, differences in the perceived credibility of the instructions, order effects and increased confusion online. Second, we investigate the differences between treatments as regards the companion measures delivered by our experiment, to see whether the differences that we identified could be explained by induced differences in secondary outcomes that might drive revealed preferences. Last, we report on companion treatments in the laboratory intended to assess the effect of the main differences in design between the online and the in-lab treatments.

4.1 Internal validity of the comparison

Our design aims to control for any treatment-specific variation in the pool of subjects by matching participants according to their registration order. Still, our sample is not large enough to guarantee a perfectly balanced sample in terms of all demographic characteristics. If any of these demographics are correlated with social preferences, then the observed differences could be driven by pool composition effects rather than the online elicitation procedure.

Table 4 provides a comparison between the two pools along all demographics available from the experiment. Out of the 12 demographic characteristics that we tested, the randomization procedure failed on one: there seem to be 7% more subjects in the laboratory sample who were not born in France.¹¹ There are no significant differences between samples in subjects' age, mothers' origin, degree level, degree level of parents, salary, student status, participation in civic organizations or religiosity.

¹¹ The table actually reports two statistically significant coefficients: one associated with the fact of not being born in France, the other associated with the fact of having a father not born in France. It turns out that these two variables are heavily related in the sample ($corr=0.51$; $p < 0.001$).

A second concern in the comparison of the two elicitation fields is a potential difference in subjects' perception about the credibility of the instructions and the payment method. Table 5 provides a summary of the self-reported assessment of the experiment stated by our subjects. Laboratory and Internet subjects report similar levels of confidence in the fact that they interact with real human partners during the experiment and will be paid at the end of the experiment as described in the instructions. We interpret these results as supportive of the internal validity of our online experimentation procedure. Further, there are also no significant differences between treatments in the care that subjects reported taking in reading the experimental instructions or in the proportion of subjects who report having participated in a similar study in the past. The only statistically significant difference that arises from this table is how calm subjects report their environment to have been when they performed the study, although the magnitude of the reported difference is small (-0.15 for Internet subjects on a 4-point scale).

Thanks to the controlled allocation of subjects across treatments, very few observable differences between the two pools arise. Moreover, the common decision platform and the overall design of the experiments have generated very few differences between subjects as regards their assessment of the credibility of the instructions. The two exceptions are the proportion of subjects who were not born in France and how calm subjects report their environment to have been when they performed the experiment. To assess the robustness of observed behavior to these dimensions, we perform separate regressions on each outcome of interest that control for all covariates (of which coefficients are omitted) and in particular these two significant differences. One last dimension that may influence our results is the possible presence of order effects. We include controls for this dimension as well. The results are reported in panel A of Table 6. We observe that the "not born in France" and "calm environment" variables have no significant impact on behavior, except for a positive and marginally significant effect of the former in the Public Good game. Similarly, the order in which games occur seems of secondary importance – as can be expected from the absence of feedback until the end of the experiment. The only exceptions concern the transfers in the Ultimatum game (order 3) and the Trust game (order 2). Importantly, we find that none of these control variables affect the estimated point differences in social preferences elicited online as compared with the laboratory.

While Table 5 shows that subjects trust the experimental instructions online and in the lab equally, we also observed in Section 3 that many more subjects appeared confused in the online risk aversion elicitation task. This raises the question of a relatively worse understanding of the instructions in this elicitation context, even though subjects reported similar levels of care in reading them. To assess the effect of this dimension, we replicate the statistical analysis of Table 6 on those subjects who showed no sign of confusion in the risk aversion task – thus using confusion in this decision problem as a proxy for confusion in the whole experiment.¹² We do not find any difference in either the significance level or even the magnitude of the relevant parameters.¹³

¹² Here we define confusion as either choosing the secure option (*i.e.* option A) in the last decision or switching back from option B to option A at least once. The results are provided as supplementary material in Table A1 – panel A.

¹³ We ran two additional robustness checks confirming the reliability of these results (tables reported as supplementary material in Table A1 – panels B and C, respectively). First, we excluded from the Internet sample all subjects who logged in to the online platform *after the target of 20 participants per experimental session had already been reached* (so that we obtained a perfectly balanced sample between laboratory and Internet subjects). We thus explored the possibility that our findings were driven by those Internet subjects who logged in to the experiment last in each session. Second, we ran the analysis on social preferences while explicitly controlling for individual levels of risk aversion in the Holt & Laury task. Contrary to Internet subjects, laboratory subjects had to incur some physical and monetary costs in order to get to the lab and play. Those costs incurred *a priori* could have made laboratory subjects relatively more willing to secure their earnings from the experiment, which could be the reason behind the higher levels of risk aversion in decision-making that we observed among laboratory subjects. This higher level of risk aversion, in turn, could have induced laboratory subjects to behave in a more conservative

4.2 Differences in underlying secondary outcomes

We now turn to a second kind of confounding factor that could challenge the inference drawn from observed preferences: the effect of the field of elicitation on secondary outcomes which may drive revealed preferences. We consider three dimensions in turn: decisions times, self-reported social preferences and the expected behavior of other subjects.

First, Shavit et al. (2001) report that participants in an Internet experiment tend to exhibit shorter decision times than classroom participants, which could, according to the literature, have a sizeable impact on behavior. Table 7 presents evidence regarding decision times in both treatments. We observe that the median time spent with the experiment among Internet subjects is 6.51 minutes lower than among laboratory subjects (Wilcoxon-Mann-Whitney test, $p < 0.0001$), with an average completion time of 34 minutes across treatments. In addition, we also observe that the variance in the time spent on the experiment is significantly higher online (two-tailed F-test, $p < 0.0001$). Notwithstanding this fact, we were surprised that none of our Internet subjects remained inactive for more than 5 minutes at any point when performing the study, which we interpret as good news for its internal validity.¹⁴

To assess the influence of this treatment effect on the preferences elicited in both fields, we include decision times in the regressions presented in Table 6. For each outcome, the decision time variable is defined as the time spent by the subject on the corresponding decision problem (from the instruction screen to the decision screen). We include it both as an additional control variable and in interaction with the online treatment so as to capture the variation in social preferences online that is induced by variations in decision times. The results are presented in panel B of Table 6. Many timing coefficients are not statistically significant. When they are, however, our estimates suggest that faster decisions are associated with more other-regarding decisions.

For instance, a one standard deviation increase in decision time is associated with a 6% decrease in the proportion of the endowment unconditionally contributed and a decrease of 0.14 points in the slope of the reaction function in the Public Good game in the lab (although only for relatively low values of the average contribution of the other group members), as well as a 8.5% decrease in the proportion of the endowment that receivers in the Ultimatum game demand online. Incidentally, it is also associated with an average decrease of 0.71 in the level of risk aversion in the Holt & Laury task (but only in the lab). These results are in line with those reported in Rubinstein (2007), Rand et al. (2012) and Lotito et al. (2013), who report that shorter decision times are associated with more pro-sociality on average.¹⁵ This evidence supports the System 1/System 2 hypothesis that shorter decision times are associated with instinctive and emotional decision processes (Kahneman 2003), which should drive subjects to behave relatively more pro-socially on average. On the other hand, the timing coefficients for the Trust game are at odds with the theory, as they indicate that higher decision times are significantly associated with an *increase* in trustworthiness.

way (*i.e.* less pro-socially) in certain games. In neither case, however, do we find any impact on the magnitude and significance of our estimates.

¹⁴ Even if online subjects do seem to play faster on average, some of them spent quite a lot of time on the experiment. One extreme case was a subject who spent more than 3 hours on the experiment without once triggering the 5-minute inactivity indicator.

¹⁵ The evidence reported in Piovesan & Wengstrom (2009) is an exception.

Focusing on our coefficients of interest, we observe that controlling for decision times has no effect on the magnitude and significance of the point differences between treatments. One exception is the difference in levels of trustworthiness exhibited by participant Bs in the Trust game, which even increases.¹⁶

Next, we explore whether the elicitation field had an impact on subjects' self-reported measures of social preferences, which could in turn have had an effect on their behavior. To do so, the final questionnaire asked subjects to answer a set of traditional survey questions about social preferences.

The result of the comparison between subject pools is reported in Table 8. We can see that no statistically significant differences arise between laboratory and Internet subjects in self-reported social preferences, except for the WVS and GSS trust question, in which roughly 9% more subjects report that "most people can be trusted" in the Internet sample ($p < 0.10$).¹⁷

Last, Table 9 provides a comparison of subjects' self reports on the expected behavior of other participants in the Public Good and Trust games between treatments. The point differences in social preferences that we identified especially strongly in the Trust game do not seem to be mediated by a modification of subjects' expectations about the behavior of others depending on the experimental context either. Indeed, the only (marginally significant) difference that arises in terms of expectations is in whether subjects report having an idea of how much the other members of their group contributed when they made their decision in the Public Good game (-9% in the Internet sample, $p < 0.10$).

4.3 The effect of the Internet-specific differences in design

As stressed in Section 2.4, our strategy in designing the experiment is to make the online and in-lab environments as similar to each other as possible, while ensuring that the in-lab conditions complied with standard practice. This led us to introduce two important differences between the two designs, so as to account for the specific constraints faced when subjects do not come to a physical laboratory to participate. First, the compensation of online subjects goes through an automated PayPal transfer, which is less immediate, and perhaps less salient, than the cash payment offered to laboratory subjects. Second, since we wanted to allow online subjects to progress within the experiment at their own pace without having to wait for others to make decisions, we implemented a sequential matching scheme between participants. Importantly, this implies that the decisions made by an online subject do not affect the outcome of his current partner, but the outcome of some future online subject. In this section, we check for the sensitivity of the observed differences in behavior between the two environments to these changes in the design, through additional laboratory experiments involving each feature in turn.

¹⁶ The change in the magnitude of these coefficients is explained by the negative correlation between the Internet treatment and average decision time, which is found to be positively and significantly associated with our measures of trust and trustworthiness.

¹⁷ These measures are very likely to be correlated with unobserved factors determining behavior in our games, and so we do not include them as control variables in the regressions.

4.3.1 Design of the robustness treatments

We ran two companion treatments in the laboratory. In the *Sequential Matching* treatment, subjects in the laboratory experiment are matched with subjects from previous sessions. In the *PayPal* treatment, participants in the laboratory experiment are paid by an automated PayPal transfer. In order to comply with the general rules of our laboratory, and avoid negative reactions both in the overall subject pool and towards our experiment, this feature of the design had to be announced at the registration stage.¹⁸ More precisely, on the webpage on which subjects confirm their willingness to participate, a preliminary screen informed them that experimental earnings would be paid through PayPal transfers. Subjects were allowed to decline participation at this stage, in which case we recorded the information available in the subject management database if provided by the subjects, *i.e.* their gender, age and student status.

Three sessions of each treatment were run in May 2013. We chose the sequence of games (as described in Section 2.2) so as to balance the overall number of sessions for each order: we ran one session of each treatment with order 2, and two sessions with order 3. Since these sessions took place after our main treatments of interest, and without an online counterpart, our control over self-selection into the elicitation field does not apply to these treatments – subjects registered on the usual first-come first-served basis for both treatments. This concern about the composition of the subject pool is reinforced by self-selection at the registration stage of the PayPal treatment, as 20% of subjects actually gave up on their registration when informed of the PayPal payment.¹⁹

Table 10 provides an overview of the demographics in the pool of subjects who participated in the SeqMatch and PayPal robustness treatments as compared with the standard laboratory one. Despite the different sample sizes (180 online as opposed to 60 in each additional treatment), we observe very few differences between the in-lab and sequential matching samples. The only significant difference that arises concerns the nationality of the father. The high refusal rate of the PayPal treatment had a greater impact on the composition of the sample, however, as PayPal subjects are on average less likely to be female, more likely to be students and religious and also younger with a lower income (although marginally significantly so).

4.3.2 Results

Figure 6 replicates the qualitative description of observed behavior of Figure 5 with the four treatments taken together. In all games, the qualitative patterns in elicited preferences remain the same. One notable feature of the figure is that the relatively low proportion of fully self-interested

¹⁸ This is unlike our Internet treatment, in which subjects were informed that the final payment would be processed through PayPal right after the introductory screen of the online platform, *i.e.* after they had already registered and logged in to participate (see section 2.3). For the present treatment, self-selection into participation due to the payment system can hardly be avoided for any payment method other than cash. Even if our laboratory usually paid subjects using PayPal (or, say, a bank transfer) we would have had to announce this in the recruitment ads, hence inducing self-selection into the overall population of potential subjects. In that sense, the selection effect that occurs in this treatment replicates the one at stake in a laboratory using PayPal as a way to dematerialize subject's payments.

¹⁹ As a comparison, it is notable that none of the subjects in the Internet treatment dropped out of the experiment at the level of the PayPal payment screen. According to the data available for this treatment, subjects who gave up on their registration at the stage of the PayPal payment explanation screen were on average 23.3 years old (as opposed to 24.6 for those who eventually participated in the experiment), 30% female (as opposed to 35%) and 56% students (as opposed to 82%).

decisions in the online treatment that we identified in Figure 5 is not replicated by either the sequential matching or the PayPal treatments. Indeed, less than 20% of subjects make no transfer in the Dictator game in the online treatment, while this proportion is more than doubled in the other three laboratory treatments (figure 6.g). For this decision, the online condition is also the only one to have its mode at an equal split of the endowment (decision made by about 25% of online subjects, as opposed to 10% or less in all other samples), while the other three treatments have a mode at zero. Similarly, less than 10% of subjects make no transfer in the Trust game in the online condition, while this proportion is again more than doubled in the other treatments (figure 6.c). This pattern is less clear-cut for the contribution decisions in the Public Good game (figure 6.a) and the transfer and threshold decisions in the Ultimatum game (figures 6.e and 6.f, respectively), but remains visible.

Another insight from Figure 6 is that the distribution of returns for online subjects in the Trust game continues to dominate the distribution of returns for all other laboratory subjects (figure 6.d). It is striking, however, that when compared with the patterns of trustworthiness exhibited in the in-lab and PayPal treatments, the pattern exhibited in the sequential matching treatment is much closer to that of the online treatment. This suggests that the point differences in trustworthiness levels that we identified between our lab and Internet conditions might be at least partly due to the sequential matching that we implemented between online subjects. This result is surprising, as one might have expected the indirect reciprocity induced by this matching procedure to weaken rather than strengthen trustworthiness.

We now turn to a more formal statistical assessment of the four treatments. We proceed in two steps. First, in panel A of Table 11 we provide estimates of the treatment effects using the same specification as in panel B of Table 6 above. We observe few differences between the baseline laboratory treatment and the sequential matching and PayPal treatments, as virtually all coefficients on those robustness treatments are insignificant. Strikingly enough, one prominent exception is the level of risk aversion, which is significantly affected by the sequential matching procedure implemented in the lab. This result is surprising, as this decision problem is the only one that does not involve interactions with other subjects.

These regression results stand as a rather weak robustness test, as they may be affected by the differences in sample size between treatments. As an additional more rigorous test of the robustness of the comparison, the two bottom panels of Table 11 provide mean comparison tests against each treatment. We compare the preferences elicited online with those elicited in each robustness treatment as a benchmark in turn. These comparisons thus inform about how well online behavior is replicated by behavior in a laboratory experiment in which subjects are, respectively, matched sequentially or paid by automated PayPal transfers. Remember that only two out of the three orders considered in our treatments of interest are implemented for the robustness treatments. We thus control for order effects in the mean comparison tests reported in the table. In line with the pattern observed in the qualitative discussion, we observe that some of the previously significant differences are no longer significant when the laboratory sessions incorporate the differences in design. Focusing on social preferences, sequential matching in the laboratory seems to replicate the higher levels of trust and trustworthiness found online in the Trust game. The higher level of donation in the Dictator game, by contrast, is robust to both changes and appears to be specific to the online elicitation field. In line with the top panel of Table 11, the risk preferences elicited online are no longer different from the ones observed in the lab, when it features either PayPal payment or sequential matching.

Overall, this exercise leads to mixed conclusions. On the one hand, the comparison confirms our main conclusion that, contrary to what is generally thought, other-regarding preferences are no less

intense online than in the laboratory. For the Dictator game, the higher level of transfers even remains strongly significant in comparison to all three laboratory situations. On the other-hand, both PayPal payment and sequential matching of subjects in the lab seem to influence revealed preferences, and account for part of the point differences we observe. This raises interesting questions, as dematerialized payment is most likely to become the standard way to remunerate subjects in online experiments, and as the indirect reciprocity involved in sequential matching could have been expected to weaken rather than strengthen social preferences. As for the purpose of this study, these results show that design choices compatible with online experimentations are not neutral on behavior, and deserve systematic experimental investigation.

5 Discussion

From the results developed in the previous sections, our main methodological conclusion is in favor of the internal validity of the preferences elicited online, thanks to the additional controls of our design. In particular, no significant difference between treatments appeared in subjects' self-reported beliefs about the accuracy of the experimental instructions. In the same vein, we found that none of our online subjects seemed to have been distracted from the experiment for more than 5 minutes (although major distractions may occur in an even shorter time-range) and that a relatively modest number of online subjects (6 out of 208) eventually dropped out of the experiment before its completion. Importantly, unlike earlier studies (*i.e.* Anderhub et al. (2001) and Shavit et al. (2001)), the dispersion of preferences that we elicit online is often statistically indistinguishable from that of the lab.

The experiment does highlight some specificities of online elicitation of behavior, though. Consistent with the above-mentioned seminal studies, we find that it is relatively more difficult to collect good quality data over the Internet, as 22 subjects on the Internet failed to select option B in the 10th decision (in which subjects had the choice between earning 20€ or 38.5€ with certainty) as compared with 5 in the laboratory. However, it should be possible to compensate for this extra noise in the data by leveraging the Internet to recruit larger samples. Finally, we find that online subjects play significantly faster on average than laboratory subjects, with sometimes a sizeable impact on behavior. Depending on the kind of experimental data, including controls for this dimension of behavior can therefore be important.

These observations speak in favor of the reliability of Internet data. The second important question this paper aimed to answer is the reliability of Internet-based inference – taking behavior in the laboratory as a benchmark. The qualitative patterns in the data unambiguously answer yes to this question, as the Internet-based experiment generates social preferences that are similar to the laboratory ones. Subjects interacting in an online setting exhibit pro-social behavior, are conditionally cooperative on average, often altruistic in the Dictator game, reveal a taste for fairness in the Ultimatum game that other subjects anticipate in the form of higher average transfers, and exhibit both trust and trustworthiness in the Trust game.

Beyond the reliability and the internal validity of social preference elicitation online, we also find that the magnitude of other-regarding behavior is not weakened by social interactions online. The amount sent in the dictator game, and the amount returned in the trust game is even significantly higher for online subjects. A more exacting assessment of the data in this regard would consist in looking statistically at the simultaneous coincidence (or difference) in social preferences elicited in

both fields. To define the null of such a test, however, one has to choose which outcomes or measures are worth considering. For instance, one could focus on one outcome variable per decision role in each game, or include all averages described in Table 3, account for decisions times as well, or even add differences in variance and the like. Instead of reporting the statistics on the joint significance of all imaginable combinations of outcomes of interest, or choosing a few particular combinations, we decided to report all results with the p -values of univariate comparisons. The Bonferroni correction for multiple comparisons can then be applied to test for joint equality of any combination of the results reported (Bland & Altman 1995). According to the correction, the threshold used to conclude on the equality of k outcomes of interest in order to replicate a Type I error equal to α is α/k . Given the strength of the statistical differences in both the trust game and the dictator game, such an exercise concludes in most instances that there is a significant difference in behavior between the two settings,²⁰ in the direction of higher other-regarding preferences online.

Given that the Internet is often viewed as the realm of anonymity (and rightly so), one might have expected the increased social distance between Internet-based subjects to drive measures of social preferences down, compared with the traditional laboratory setting. For instance, Hoffman et al. (1996) show that subjects tend to decrease the amount of their transfers in the Dictator game when social distance (*i.e.* isolation) increases and Glaeser et al. (2000) report that measures of trust and trustworthiness tend to increase with the level of demographic similarity between both players. As regards social distance theory, two alternative conclusions can be drawn from this observation. It challenges either the generally acknowledged greater social distance that prevails on the Internet (Fiedler et al. 2011), or the prediction of social distance theory per se. Our data cannot distinguish between these two views of our results.

A tentative alternative explanation can be found in the nature of many of the social and economic interactions in which individuals tend to engage online, which they may bring to the experiment through its contextual implementation. As the Internet is an environment in which it is difficult to enforce contracts, trust and trustworthiness are likely to be major devices through which to secure online transactions and build a reputation for oneself (Greif 2006). So perhaps the strong anonymity that prevails in Internet-based interactions does not come at the expense of social preferences.²¹ The prominent role of trust and trustworthiness in Internet-based economic transactions has already been demonstrated in the case of a popular online auction site (Resnick et al. 2006). In a similar fashion, the drastic reduction in communication and coordination costs brought about by the Internet has made it easier for individuals to behave altruistically towards one another, as exemplified by the impressive growth of question-driven online message boards and customer review systems.

In a recent paper, Hoffman and Morgan (2011) explored the hypothesis that selection pressures resulting from high competition, low entry and exit barriers and agents' anonymity in online business environments should drive individuals with strong social preferences out of those markets. They got

²⁰ The exact p -value on the test of mean equality in transfers in the dictator game from Table 3 is $7.39e-7$, which drives rejection even if one accounts for more than 1000 outcomes. If we instead focus separately on positive transfers and conditional transfers, *i.e.* restricting to positive contributions only, the p -value of the difference in contributions in the dictator game is 0.0003 leading to more mixed conclusions (in the trust game, the p -value on the share of positive returns is 0.015, it is 0.0212 for the comparison in mean amounts returned if positive). For instance, the equality in social preferences between the in-lab and online treatments is rejected at the 1% level if we consider that each game yields one outcome of interest per decision role (*i.e.* $k=6$, adjusted threshold=0.0017), or if we consider each variable reported in Table 3 as one outcome of interest (*i.e.* $k=14$; adjusted threshold =0.0007). The conclusion is reversed if the variance of outcome behavior (14 outcomes), as well as the beliefs over the experiment (5) and the self reported measures of trust (5) are accounted for ($k=38$; adjusted threshold =0.00026).

²¹ The lack of an "institutional" way of securing social and economic interactions over the Internet is often invoked as a reason why many Internet users who value their anonymity online are nonetheless willing to stick to and invest in a unique online identity or pseudonym.

professionals from the Internet domain trading and online adult entertainment industries to perform a series of social preference experiments and compared the results to those obtained from a population of undergraduate students. Contrary to what they initially expected, they found that Internet business people are significantly more altruistic, more trusting, more trustworthy and less likely to lie. They interpreted these findings as support for the idea that social preferences are rewarded in the Internet environment, where they help to smooth interactions and are thus beneficial in the long run. Again, our study was not designed to test this explanation against any of a possible set of alternative hypotheses. Future studies should dig into the precise nature of this “Internet effect” that we have found.

6 Conclusion

The Internet is becoming increasingly attractive to experimenters, both as a *medium* through which to target larger and more diverse samples with reduced administrative and financial costs, and as a *field* of social science research in its own right. In this paper, we report on a randomized experiment eliciting social preferences and risk aversion both online and in the laboratory based on the same, original, Internet-based platform. To provide a testbed comparison of social experimentations online, our platform seeks to control for most of the dimensions commonly highlighted as possibly challenging their internal validity, including self-sorting, differences in response times, concentration and distraction, or differences in experimental instructions and payment methods, together with their credibility.

This testbed comparison shows that online elicitation of preferences is internally valid, according to the additional controls of our design. In particular we find that the qualitative patterns of preferences elicited in the lab are often indistinguishable from those elicited online, whether in terms of treatment effects, point differences or behavioral variance. We do find, however, that it is relatively more difficult to collect good quality data over the Internet – as shown by the increase in the number of inconsistencies in the risk aversion elicitation task. However, it should be possible to compensate for this extra noise in the data by leveraging the Internet to recruit larger samples. Last, we obtain some interesting counterintuitive results as regards social preferences exhibited online. Irrespective of whether the point differences are statistically significant or not, our results indicate that when compared to subjects allocated to the laboratory condition, other-regarding behavior from subjects in the Internet condition is never weaker – sometimes stronger. Those results are at odds with what social distance theory and common wisdom predict, given that the Internet is often characterized as an environment where anonymity is more stringent. As the online environment arguably relies more on trust to achieve trade and contract enforcement, we suggest that such habits may outperform the effect of increased social distance.

These findings are important to the growing community of researchers interested in using the Internet to run large-scale social experiments online and relating their results to the established laboratory literature. Provided that enough care is taken over specific aspects of the design, Internet-based experimental inference should be considered reliable, and the results obtained from online experiments can be compared to those obtained in the lab. These results are also potentially important for social scientists wishing to use social experiments to research the Internet as a field.

Our study raises several unanswered questions. First, we apply our methodology to the elicitation of social preferences – because there were strong reasons to doubt the parallelism between the two

fields – but many other dimensions of preferences or strategic decision-making could vary between the two environments. Second, while our design appears to be adequate to guarantee the internal validity of the preferences elicited over the Internet, our experiment was not designed to differentiate the specific dimensions that were most crucial to achieving this outcome. This is an important issue to investigate in the future, as our results have shown that some design choices compatible with online experimentations are not neutral to behavior. Last, insofar as we do observe some differences in revealed social preferences between the two elicitation fields, we are unable to conclude which of the two measures is closer to actual economic behavior. Actual differences in revealed preferences depending on the field of decision elicitation, and which field scholars should trust more, warrants a more systematic investigation which we leave open for future research.

References

- Akerlof, G.A., 1997. Social Distance and Social Decisions. *Econometrica*, 65(5), pp.1005–1027.
- Amir, O., Rand, D.G. & Gal, Y.K., 2012. Economic Games on the Internet: The Effect of \$1 Stakes. *PLoS ONE*, 7(2), pp.1–4.
- Anderhub, V., Müller, R. & Schmidt, C., 2001. Design and evaluation of an economic experiment via the Internet. *Journal of Economic Behavior & Organization*, 46(2), pp.227–247.
- Bainbridge, W.S., 2007. The Scientific Research Potential of Virtual Worlds. *Science*, 317(5837), pp.472–476.
- Bland, J.M. & Altman, D.G., 1995. Multiple significance tests: the Bonferroni method. *BMJ: British Medical Journal*, 310(6973), p.170.
- Charness, G., Haruvy, E. & Sonsino, D., 2007. Social distance and reciprocity: An Internet experiment. *Journal of Economic Behavior & Organization*, 63(1), pp.88–103.
- Chesney, T., Chuah, S.-H. & Hoffmann, R., 2009. Virtual world experimentation: An exploratory study. *Journal of Economic Behavior & Organization*, 72(1), pp.618–635.
- Cooper, D.J. & Saral, K.J., 2013. Entrepreneurship and team participation: An experimental study. *European Economic Review*, 59, pp.126–140.
- Dohmen, T. et al., 2011. Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences. *Journal of the European Economic Association*, 9(3), pp.522–550.
- Eckel, C.C. & Wilson, R.K., 2006. Internet cautions: Experimental games with internet partners. *Experimental Economics*, 9(1), pp.53–66.
- 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.
- Fiedler, M. & Haruvy, E., 2009. The lab versus the virtual lab and virtual field—An experimental investigation of trust games with communication. *Journal of Economic Behavior & Organization*, 72(2), pp.716–724.
- Fiedler, M., Haruvy, E. & Li, S.X., 2011. Social distance in a virtual world experiment. *Games and Economic Behavior*, 72(2), pp.400–426.
- Fischbacher, U., Gächter, S. & Fehr, E., 2001. Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3), pp.397–404.
- Glaeser, E.L. et al., 2000. Measuring Trust. *The Quarterly Journal of Economics*, 115(3), pp.811–846.
- Greif, A., 2006. *Institutions And The Path to the Modern Economy: Lessons from Medieval Trade*, Cambridge University Press.
- Greiner, B., 2004. *An Online Recruitment System for Economic Experiments*, University Library of Munich, Germany.

- Henrich, J., Heine, S.J. & Norenzayan, A., 2010. The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), pp.61–83.
- Hoffman, E., McCabe, K. & Vernon L. Smith, 1996. Social Distance and Other-Regarding Behavior in Dictator Games. *The American Economic Review*, 86(3), pp.653–660.
- Hoffman, M. & Morgan, J., 2011. Who's Naughty? Who's Nice? Social Preferences in Online Industries. *UC Berkeley Working Paper*.
- Holt, C.A. & Laury, S.K., 2002. Risk Aversion and Incentive Effects. *The American Economic Review*, 92(5), pp.1644–1655.
- Horton, J.J., Rand, D.G. & Zeckhauser, R.J., 2011. The online laboratory: conducting experiments in a real labor market. *Experimental Economics*, 14(3), pp.399–425.
- Kahneman, D., 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review*, 93(5), pp.1449–1475.
- Lotito, G., Migheli, M. & Ortona, G., 2013. Is cooperation instinctive? Evidence from the response times in a public goods game. *Journal of Bioeconomics*, 15(2), pp.123–133.
- Piovesan, M. & Wengström, E., 2009. Fast or fair? A study of response times. *Economics Letters*, 105(2), pp.193–196.
- Rand, D.G., Greene, J.D. & Nowak, M.A., 2012. Spontaneous giving and calculated greed. *Nature*, 489(7416), pp.427–430.
- Resnick, P. et al., 2006. The value of reputation on eBay: A controlled experiment. *Experimental Economics*, 9(2), pp.79–101.
- Rubinstein, A., 2007. Instinctive and Cognitive Reasoning: A Study of Response Times. *The Economic Journal*, 117(523), pp.1243–1259.
- Selten, Reinhard, 1967. Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopol experiments. In Sauermann, H, ed. *Beitrage zur Experimentellen Wirtschaftsforschung*. Tübingen: J.C.B. Mohr, pp. 136–168.
- Shavit, T., Sonsino, D. & Benzion, U., 2001. A comparative study of lotteries-evaluation in class and on the Web. *Journal of Economic Psychology*, 22(4), pp.483–491.

Tables and Figures

Table 1. In-lab versus online based experiment: overview of experimental results

Paper	Type of experiment	Subject pool	Random allocation of subjects	Main results
Anderhub et al. (2001)	Individual level consumption/saving decisions	47 in lab 50 online	NO	(i) similar economic behavior on average (ii) higher behavioral variance online (iii) shorter decision times online
Shavit et al. (2001)	Individual lotteries evaluation decisions	65 in classroom 70 online	NO	(i) lower risk aversion online (ii) higher behavioral variance online
Charness et al. (2007)	Lost wallet game	178 in classroom 124 online	NO	Very little difference in average economic behavior
Fiedler and Haruvy (2009)	Trust game with pre-play communication	136 in lab 216 online	NA	Lower levels of trust and trustworthiness online
Chesney et al. (2009)	Dictator game, Ultimatum game, Public Good game, Minimum Effort game, Guessing game	Respectively 30, 64, 32, 31 and 31 online	NA	Behavioral results qualitatively in line with previous laboratory based experiments
Horton et al. (2011)	Watershed experiment, Religiously primed and unprimed versions of the Prisoner's Dilemma	Respectively 213, 189 and 113 online	NA	Behavioral results qualitatively in line with previous laboratory based experiments
Amir et al. (2012)	Public Good game, Dictator game, Ultimatum game, Trust game	189 per game online	NA	Behavioral results qualitatively in line with previous laboratory based experiments

Figure 1. The description screen of the Trust game

Section 3/4 - Description

In this section, pairs of 2 participants (yourself and one other participant) are randomly matched. Each pair has a participant A and a participant B. Your role (participant A or B) is indicated at the moment when you have to make your decision.

Remember: The participant who belongs to your pair in this section is different from those you encounter in the other sections of the study.

At the beginning of this section, each of the two participants receives \$10.

Participant A must decide how many dollars out of his \$10 will be transferred to participant B. The amount transferred by participant A is multiplied by 3. Thus, for each dollar that participant A transfers, participant B receives \$3.

Participant B must then decide how much to return to participant A. Before finding out how much participant A is going to transfer, participant B must decide how much to return for each possible transfer (all integers between 1 and 10).

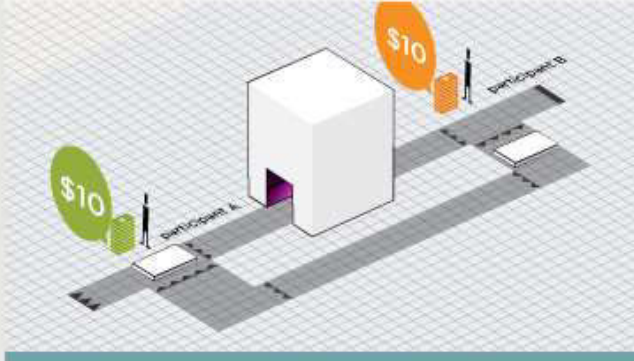
The earnings in dollars of participant A at the end of this section are given by:

$$10 - (\text{amount in dollars transferred to participant B}) + (\text{amount in dollars returned by participant B})$$

The earnings in dollars of participant B for this section are given by:

$$10 + (\text{amount in dollars transferred by participant A}) \times 3 - (\text{amount in dollars returned to participant A})$$

⇒ The next screen gives examples...



[Previous](#) [Next](#)

Figure 2. Flash animation for the Public Good game





Figure 3. Earnings calculators



(a) Public Good game



(b) Trust game

Figure 4. Decision screen for the Public Good game

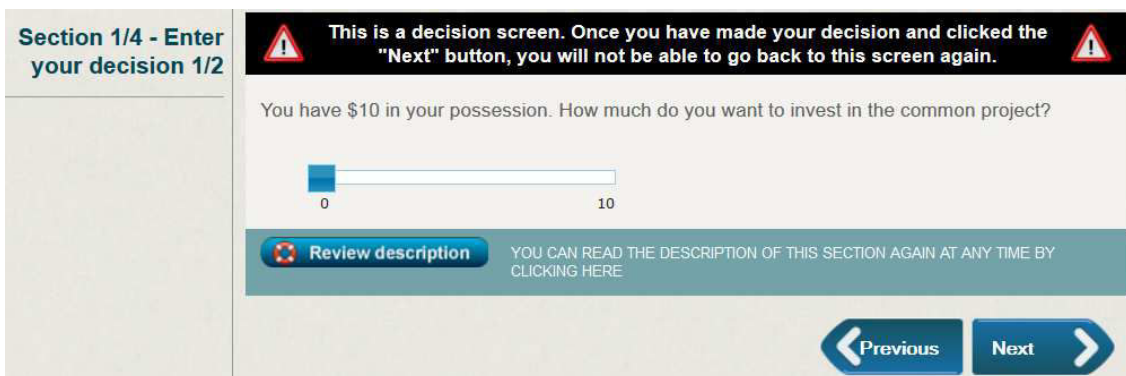


Table 2.1. Summary of the design: common procedures between treatments


Decision problems	Decisions elicited from participant		Games ordering			Sequence of screens
	A	B	1	2	3	
1. PUBLIC GOOD GAME	(i) unconditional contribution (ii) conditional contribution (strategy method)		1	1	1	- Description (text + animation) - Illustrative examples - Earnings calculator - Decision screen unconditional - Decision screen conditional - Beliefs elicitation
<i>Elicitation of beliefs</i>	(i) normative view on how much people should contribute (ii) idea about contributions of others at time of decision (iii) estimation of contributions of others at time of decision					
2. DICTATOR GAME	Transfer	<i>None</i>	2	4	3	- Description (text + animation) - Decision screen
3. ULTIMATUM GAME	Transfer	Minimum acceptable offer	3	3	2	- Description (text + animation) - Decision screen
4. TRUST GAME	Transfer	Amount returned (strategy method)	4	2	4	- Description (text + animation) - Illustrative examples - Earnings calculator - Decision screen - Beliefs elicitation
<i>Elicitation of beliefs</i>	(i) idea about return at time of decision (ii) estimation of return at time of decision	<i>None</i>				
5. HOLT & LAURY LOTTERIES	Choice over 10 lottery pairs		5	5	5	- Description (text + illustrative table) - Decision screen
Social values survey	Cooperation, altruism, fairness, trust (WVS), general trust, trust in strangers, risk aversion (see table 8)					
Debriefing Questionnaire	(i) demographic control variables (see table 4) (ii) beliefs over the experiment (see table 5)					

Table 2.2. Summary of the design: differences in implementation between treatments

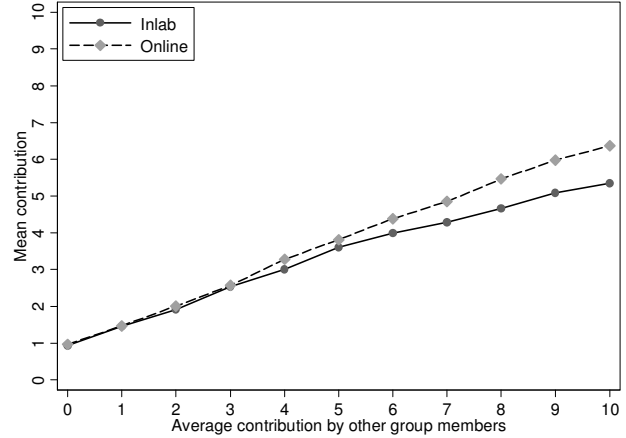
	Matching	Payment	Participation slot	Overbooked subjects
Inlab	Simultaneous	Cash	At time	Refused
Online	Sequential	Automated PayPal transfer	Any time during the half-day spanning the slot	Identified in the data and allowed to participate

Figure 5. Behavior in the decision problems between treatments


(a) Distribution of unconditional contributions in the Public Good game




(b) Conditional contributions in the Public Good game




(c) Distribution of transfers in the Trust game




(d) Amounts returned in the Trust game




(e) Distribution of transfers in the Ultimatum game



(f) Distribution of minimum acceptable offers in the Ultimatum game



(g) Distribution of transfers in the Dictator game



(h) Risk aversion levels in the Holt&Laury task




Table 3. Descriptive statistics

Variable	Nb Of Obs.		Mean behaviors			Standard deviation		
	Inlab	Online	Inlab	Online	<i>p</i> -value	Inlab	Online	<i>p</i> -value
<i>Public Good Game</i>								
Contribution	180	202	3.64	3.89	0.2028	3.06	2.73	0.1202
Mean conditional contributions	180	202	3.35	3.74	0.0394**	1.99	2.10	0.4567
Slope against low contributions others	180	202	0.53	0.57	0.6866	0.56	0.52	0.2870
Slope against high contributions others	180	202	0.35	0.51	0.0437**	0.73	0.61	0.0178**
<i>Dictator Game</i>								
Positive transfer	90	102	0.58	0.81	0.0004***	0.50	0.39	0.0203**
Transfer	90	102	1.62	3.36	0.0000***	1.88	2.53	0.0048***
<i>Ultimatum Bargaining Game</i>								
Transfer	90	102	4.28	4.72	0.4133	4.28	4.72	0.7469
Transfer threshold	90	100	3.00	3.69	0.0556*	1.90	2.14	0.2582
<i>Trust Game</i>								
Amount sent	90	102	3.54	4.45	0.0193**	3.32	3.01	0.3360
Mean amounts returned	90	100	3.85	6.29	0.0001***	3.72	4.33	0.1473
Slope against low amounts sent	90	100	0.67	1.10	0.0007***	0.72	0.82	0.2397
Slope against high amounts sent	90	100	0.71	1.20	0.0016***	0.91	0.98	0.4624
<i>Holt&Laury lottery choices</i>								
Nb of safe choices	180	202	6.76	6.15	0.0021***	1.78	2.03	0.0771*
Nb of safe choices w/o confused	164	152	6.80	6.18	0.0075***	1.70	2.01	0.0345**

Notes: *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *p*-values are from Wilcoxon-Mann-Whitney tests (for differences in distributions) and two-sided variance comparison tests (for differences in variances), respectively. *Public Good game*: *Contribution* = unconditional contribution to the common project; *Mean conditional contributions* = mean of conditional contributions to the common project; *Slope against low contributions others* = slope of the reaction function for average contributions of other group members from 0 to 5; *Slope against high contributions others* = slope of the reaction function for average contributions of other group members from 6 to 10. *Dictator game*: *Positive transfer* = transfer in the Dictator game is positive; *Transfer* = transfer in the Dictator game. *Ultimatum game*: *Transfer* = transfer in the Ultimatum game; *Transfer threshold* = minimum acceptable offer in the Ultimatum game. *Trust game*: *Amount sent* = amount transferred in the Trust game; *Mean amounts returned* = mean of the amounts returned to participant A; *Slope against low amounts sent* = slope of the reaction function for amounts transferred by participant A from 1 to 5; *Slope against high amounts sent* = slope of reaction function for amounts transferred by participant A from 6 to 10. *Holt&Laury lottery choices*: *Nb of safe choices* = number of times (out of 10) the subject chose the secure option (*i.e.* option A); *Nb of safe choices w/o confused* = number of times (out of 10) the subject chose the secure option (*i.e.* option A) excluding the sub-sample of inconsistent subjects, *i.e.* all subjects who either chose the secure option (*i.e.* option A) in the last decision or switched back from option B to option A at least once.

Table 4. Demographic characteristics between treatments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Age	Female	Not born in France			Highest degree completed			Salary	Student	Participates in civic organization	Religious Person
			Subject	Father	Mother	Subject	Father	Mother				
Online	0.0436	0.0564	-0.0706*	-0.103**	-0.0237	-0.192	-0.371	-0.200	-0.0034	-0.0151	0.0717	0.0272
(<i>p</i> -value)	(0.969)	(0.269)	(0.0865)	(0.0423)	(0.642)	(0.213)	(0.169)	(0.431)	(0.977)	(0.760)	(0.104)	(0.548)
N	382	382	382	382	382	381	262	266	372	382	382	382
R ²	0.000	0.003	0.008	0.011	0.001	0.004	0.007	0.002	0.000	0.000	0.007	0.001

Notes: OLS estimates with baseline=Inlab. *p*-values are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Constants not reported.

Table 5. Beliefs over the experiment

	(1)	(2)	(3)	(4)	(5)
	Believes others are human subjects	Believes final payment will be proceeded	Has read the instructions carefully	The environment was calm	Has already participated in similar study
Online	0.0655	-0.0408	-0.0198	-0.1510**	-0.0107
(<i>p</i> -value)	(0.579)	(0.662)	(0.788)	(0.021)	(0.832)
N	265	271	382	382	382
R ²	0.001	0.001	0.000	0.014	0.000

Notes: OLS estimates with baseline=Inlab. *p*-values are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Constants not reported.

Table 6. Regression analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	<i>Public Good</i>				<i>Dictator</i>	<i>Ultimatum</i>		<i>Trust</i>				<i>Holt&Laury lotteries</i>	
	Contribution	Mean conditional contributions	Slope against low	Slope against high	Transfer	Transfer	Transfer threshold	Amount sent	Mean amounts returned	Slope against low	Slope against high	Nb safe choices	Nb safe choices <i>w/o confused</i>
Panel A: Includes controls for (i) demographic characteristics (ii) beliefs over the experiment and (iii) games ordering controls													
Online	0.187 (0.60898)	0.0938 (0.72433)	0.0636 (0.36483)	0.112 (0.18484)	1.945*** (0.00000)	0.609* (0.05607)	0.653 (0.11764)	1.102* (0.05041)	1.996*** (0.00902)	0.337** (0.02273)	0.292 (0.10589)	-0.683*** (0.00688)	-0.822*** (0.00248)
Not born in France	0.816* (0.08282)	0.599* (0.07961)	-0.0137 (0.87859)	-0.126 (0.24456)	0.644 (0.16216)	0.564 (0.13873)	-0.0129 (0.98214)	0.847 (0.20748)	0.543 (0.60463)	0.138 (0.49630)	-0.0621 (0.80348)	-0.0731 (0.81990)	-0.0701 (0.84369)
Calm environment	0.184 (0.55870)	-0.0119 (0.95841)	-0.00846 (0.88832)	-0.0215 (0.76675)	0.0868 (0.77946)	-0.0611 (0.81170)	0.158 (0.67515)	0.204 (0.65192)	-0.329 (0.63028)	-0.0747 (0.57342)	-0.165 (0.31130)	0.266 (0.21746)	0.0455 (0.84918)
Games order 2	0.199 (0.64939)	0.133 (0.67594)	-0.0389 (0.64324)	-0.0157 (0.87654)	0.595 (0.19400)	0.384 (0.31031)	0.496 (0.31252)	1.382** (0.03979)	-0.405 (0.64863)	-0.0997 (0.56335)	-0.0674 (0.75003)	0.155 (0.60526)	-0.202 (0.53396)
Games order 3	0.551 (0.21740)	0.222 (0.49261)	0.00994 (0.90736)	0.0269 (0.79319)	0.170 (0.70648)	1.031*** (0.00667)	-0.134 (0.79230)	0.308 (0.64088)	-1.419 (0.12614)	-0.204 (0.25645)	-0.262 (0.23456)	0.00726 (0.98103)	0.0255 (0.93672)
Constant	0.0904 (0.96118)	3.111** (0.02178)	0.508 (0.15391)	0.189 (0.65779)	-0.610 (0.74134)	3.772** (0.01474)	1.372 (0.55247)	-0.554 (0.83707)	7.451* (0.07734)	1.040 (0.20231)	2.580** (0.01085)	4.515*** (0.00046)	6.312*** (0.00002)
R ²	0.103	0.085	0.059	0.087	0.343	0.162	0.114	0.200	0.205	0.186	0.219	0.090	0.119
Panel B: Same as Panel A and (iv) game specific decision times													
Online	-0.516 (0.28799)	-0.200 (0.55616)	-0.106 (0.23494)	0.144 (0.18396)	1.847*** (0.00001)	0.688** (0.03475)	0.609 (0.14384)	1.140* (0.06146)	2.596*** (0.00170)	0.472*** (0.00300)	0.410** (0.03868)	-0.636** (0.01050)	-0.733*** (0.00634)
Game specific timing	-0.581* (0.07475)	0.00596 (0.98117)	-0.137** (0.03887)	0.113 (0.16210)	0.339 (0.46458)	0.508 (0.15725)	-0.0711 (0.81799)	0.575 (0.32477)	1.097* (0.07731)	0.248** (0.03857)	0.216 (0.15055)	-0.705*** (0.00897)	-0.589** (0.03323)
Game specific timing x online	0.0913 (0.85310)	-0.541 (0.11845)	-0.0137 (0.87945)	-0.197* (0.07612)	-0.965* (0.07045)	-0.504 (0.20893)	-0.778* (0.05103)	-0.732 (0.31528)	-0.196 (0.79790)	-0.0705 (0.63239)	-0.106 (0.56624)	0.616** (0.04433)	0.458 (0.14946)
Constant	0.339 (0.85516)	2.719** (0.04541)	0.582 (0.10119)	0.0629 (0.88428)	-0.119 (0.94873)	3.690** (0.01783)	0.296 (0.89629)	-0.973 (0.72344)	6.676 (0.11092)	0.849 (0.29144)	2.388** (0.01963)	4.963*** (0.00012)	6.810*** (0.00000)
R ²	0.121	0.108	0.099	0.097	0.378	0.177	0.182	0.208	0.256	0.247	0.241	0.115	0.141
N	257	257	257	257	131	131	126	131	126	126	126	257	207

Notes: OLS estimates with baseline=Inlab. *p*-values are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Demographic controls are all variables from table 4. Beliefs over the experiment controls are all variables from table 5. Game specific timing variables are standardized. *Public Good Game*: Contribution = unconditional contribution to the common project; Mean conditional contributions = mean of conditional contributions to the common project; Slope against low = slope of the reaction function for average contributions of other group members from 0 to 5; Slope against high = slope of the reaction function for average contributions of other group members from 6 to 10. *Dictator game*: Transfer = transfer in the Dictator game. *Ultimatum game*: Transfer = transfer in the Ultimatum game; Transfer threshold = minimum acceptable offer in the Ultimatum game. *Trust game*: Amount sent = amount transferred in the Trust game; Mean amounts returned = mean of the amounts returned to participant A; Slope against low = slope of the reaction function for amounts transferred by participant A from 1 to 5; Slope against high = slope of the reaction function for amounts transferred by participant A from 6 to 10. *Holt&Laury lotteries*: Nb safe choices = number of times (out of 10) the subject chose the secure option (i.e. option A); Nb safe choices w/o confused = number of times (out of 10) the subject chose the secure option (i.e. option A) excluding the sub-sample of confused subjects, i.e. all subjects who either chose the secure option (i.e. option A) in the last decision or switched back from option B to option A at least once.

Table 7. Difference in median/variance of time spent on the experiment

Number of Observations		Median time			Standard Deviation		
Inlab	Online	Inlab	Online	Diff.	Inlab	Online	Diff.
120	154	35.01	28.50	6.51***	7.77	17.52	- 9.74***
				$p < 0.0001$			$p < 0.0001$

Notes: p -values are from a Wilcoxon-Mann-Whitney test (for difference in distributions) and two-sided variance comparison tests (for difference in variances), respectively. *, ** and *** denote statistical significance at the 10, 5 and 1% levels.

Table 8. Self-reported social preferences between treatments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Cooperation	Altruism	Fairness	Trust (WVS)	General trust	Trust in strangers	Riskaver
Online	0.457	0.148	-0.235	0.0887*	-0.0477	-0.0551	0.300
(p -value)	(0.117)	(0.474)	(0.271)	(0.0676)	(0.477)	(0.447)	(0.247)
N	366	376	372	352	370	372	271
R ²	0.007	0.001	0.003	0.010	0.001	0.002	0.005

Notes: OLS estimates of column variables on the online dummy (the baseline is inlab subjects, constants are not reported). *, ** and *** denote statistical significance at the 10, 5 and 1% levels. *Cooperation* = whether subjects consider it justifiable to free-ride on public social allowances; *Altruism* = whether subjects think that people are mostly looking out for themselves as opposed to trying to help each other; *Fairness* = whether subjects think that people would try to take advantage of them if they got a chance as opposed to trying to be fair; *Trust (WVS)* = whether subjects think that most people can be trusted or that one needs to be very careful when dealing with people; *General trust* = subjects' level of general trust in people; *Trust in strangers* = how much trusting subjects are of people they just met; *Riskaver* = whether subjects generally see themselves as fully prepared to take risks or as trying to avoid them.

Table 9: Beliefs about other subjects' decisions by treatment

	(1)	(2)	(3)	(4)	(5)
	How much others should contribute	Idea about how much others will contribute	Estimation of how much others will contribute	Idea about how much the responder will return	Estimation of how much the responder will return
Online	-0.450	-0.0910*	0.202	-0.0719	0.0737
(<i>p</i> -value)	(0.147)	(0.0538)	(0.496)	(0.260)	(0.584)
N	381	382	266	192	116
R ²	0.006	0.010	0.002	0.007	0.003

Notes: OLS estimates with baseline=Inlab. *p*-values are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Constants not reported. (1) is how much subjects think people *should* contribute to the common project in the Public Good game; (2) is whether subjects had an idea of how much the other subjects in their group *would actually* contribute to the common project when they made their decision; (3) is conditional on (2), how much subjects thought the other subjects in their group would contribute on average when they made their decisions. (4) is whether subjects in the role of senders in the Trust game had an idea of how much the responder would return to them when they made their decision; (5) is conditional on (4), proportion of the amount sent that trustors anticipated would be returned to them by the trustee when they made their decision.

Table 10: Demographic characteristics between the in-lab, sequential matching and PayPal treatments

	(1)	(2)	(3) (4) (5)			(6) (7) (8)			(9)	(10)	(11)	(12)
	Age	Female	Not born in France			Highest degree completed			Salary	Student	Participates in civic organization	Religious Person
			Subject	Father	Mother	Subject	Father	Mother				
SeqMatch	2.628	0.0389	-0.0556	-0.178**	-0.0944	-0.246	-0.316	0.0796	0.0679	-0.0889	-0.00556	0.0500
(<i>p</i> -value)	(0.110)	(0.600)	(0.382)	(0.0165)	(0.202)	(0.269)	(0.369)	(0.814)	(0.690)	(0.205)	(0.926)	(0.460)
PayPal	-2.824*	-0.178**	0.0444	0.0722	0.106	0.338	-0.203	0.119	-0.299*	0.178**	-0.0222	0.183***
(<i>p</i> -value)	(0.0859)	(0.0169)	(0.485)	(0.328)	(0.154)	(0.127)	(0.559)	(0.721)	(0.0843)	(0.0115)	(0.711)	(0.00708)
N	382	382	382	382	382	381	262	266	372	382	382	382
R ²	0.000	0.003	0.008	0.011	0.001	0.004	0.007	0.002	0.000	0.000	0.007	0.001

Notes: OLS estimates with baseline=Inlab. *p*-values are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Constants not reported.

Figure 6. Behavior in the decision problems between treatments (including the SeqMatch and PayPal treatments)

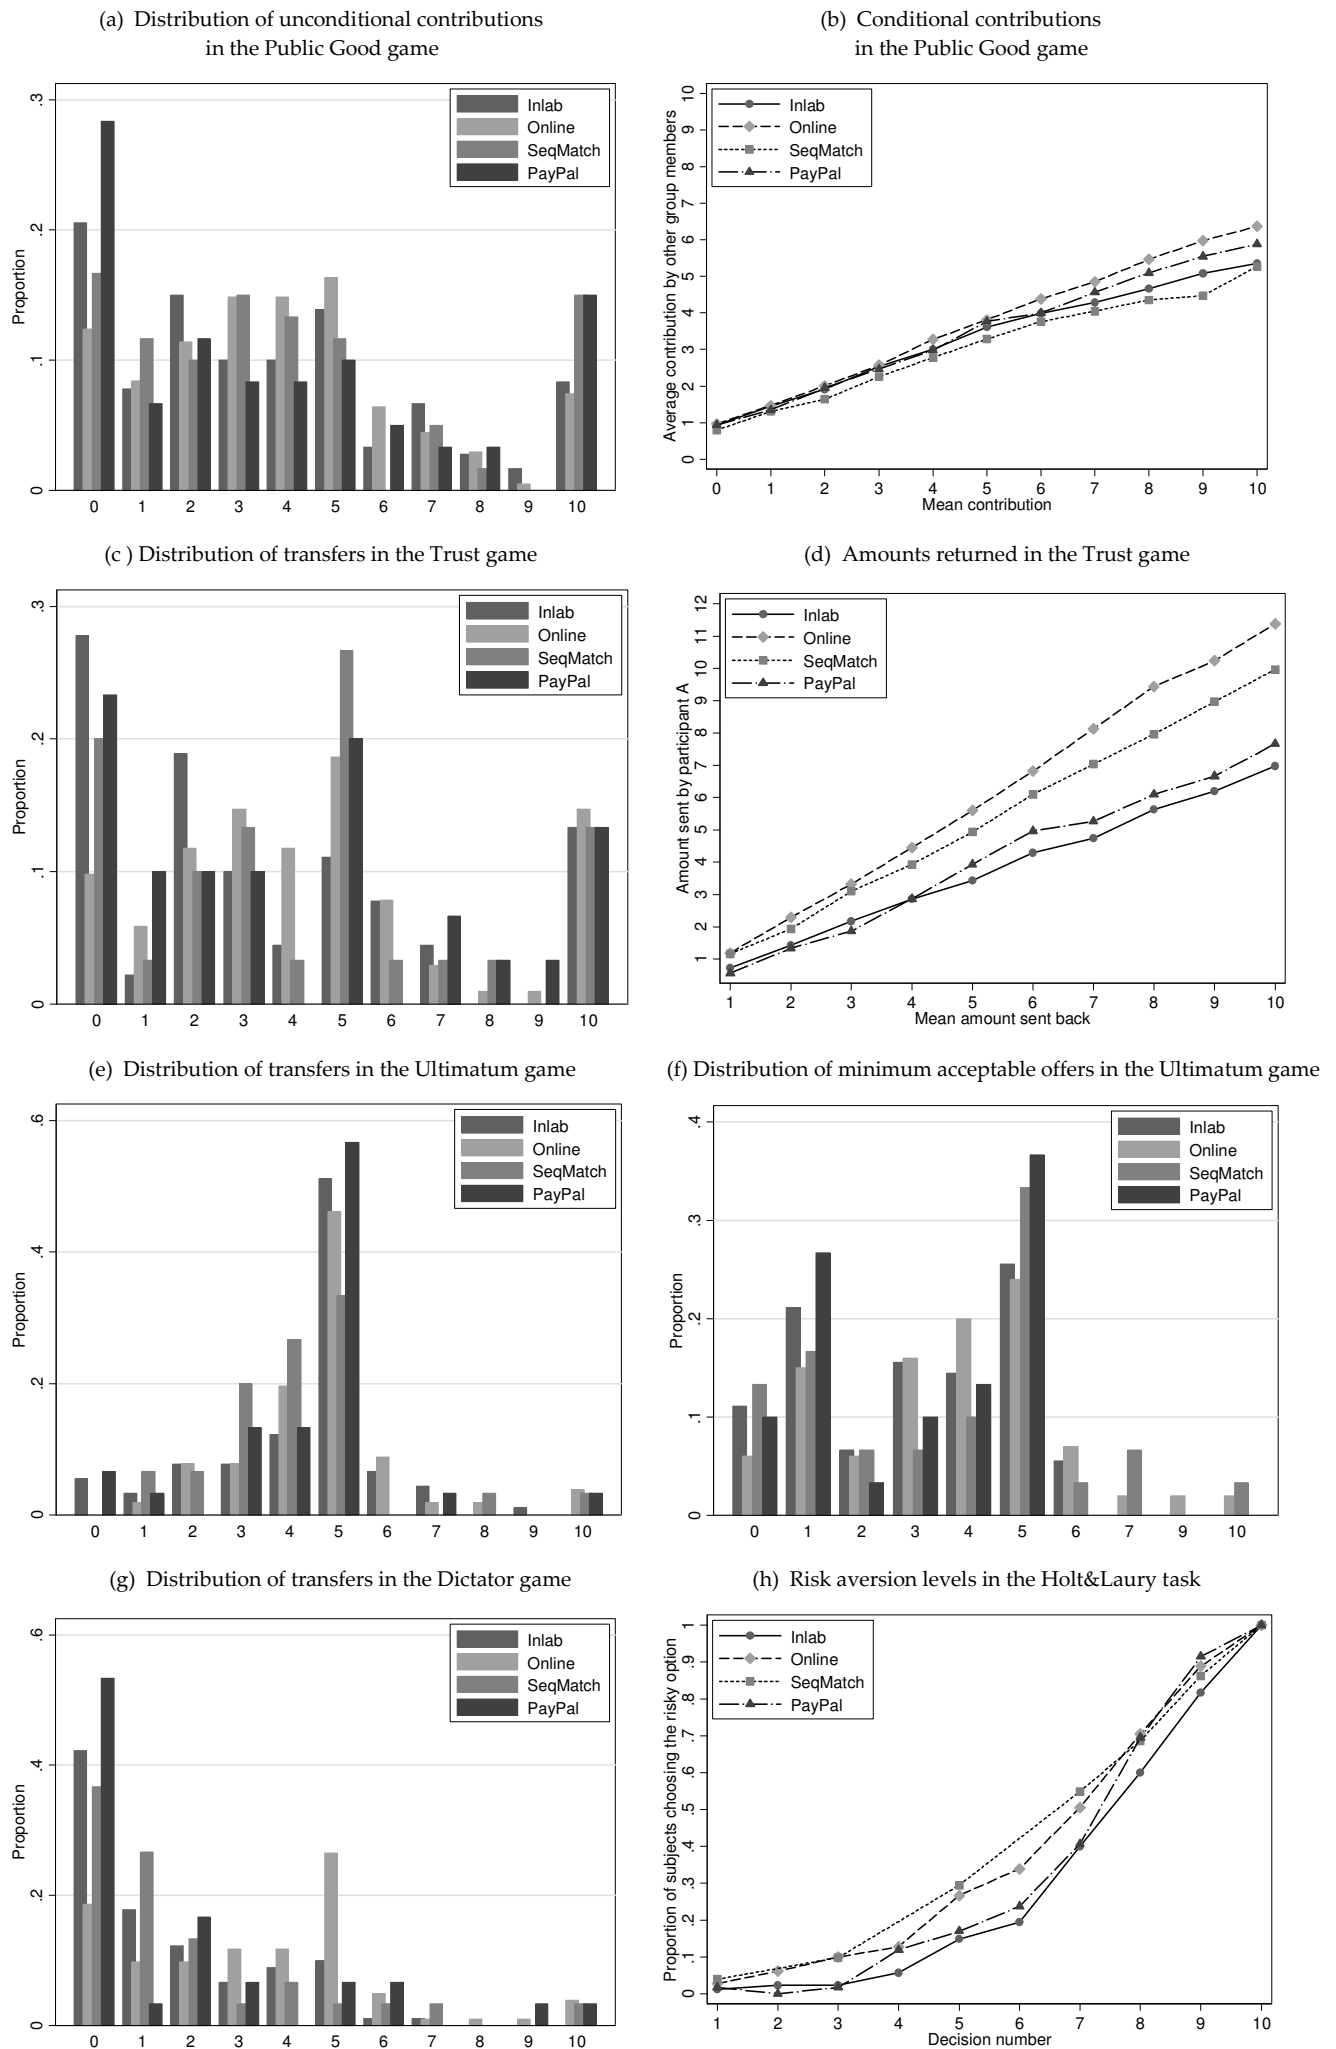


Table 11. The effect of sequential matching and PayPal payment on behavior

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	<i>Public Good</i>				<i>Dictator</i>	<i>Ultimatum</i>		<i>Trust</i>				<i>Holt&Laury lotteries</i>	
	Contribution	Mean conditional contributions	Slope against low	Slope against high	Transfer	Transfer	Transfer threshold	Amount sent	Mean amounts returned	Slope against low	Slope against high	Nb safe choices	Nb safe choices <i>w/o</i> confused
Panel A: All treatments pooled – Baseline=Inlab treatment													
Online	-0.287 (0.54645)	-0.160 (0.63026)	-0.0395 (0.61771)	0.0864 (0.38877)	1.867*** (0.00000)	0.594* (0.06187)	0.631 (0.10752)	1.158* (0.05352)	2.808*** (0.00019)	0.479*** (0.00094)	0.470** (0.01078)	-0.612** (0.01405)	-0.682** (0.01016)
SeqMatch	-0.560 (0.29866)	-0.394 (0.31632)	-0.0551 (0.55626)	0.0267 (0.82177)	0.516 (0.35814)	0.0158 (0.97291)	0.270 (0.59898)	0.580 (0.46843)	1.489 (0.11413)	0.205 (0.25960)	0.222 (0.34017)	-0.695** (0.03586)	-0.929*** (0.00765)
PayPal	-0.401 (0.47030)	0.128 (0.75264)	-0.000327 (0.99730)	-0.00278 (0.98180)	0.524 (0.33252)	0.0111 (0.98042)	0.130 (0.80956)	0.637 (0.42938)	1.324 (0.17955)	0.322* (0.09202)	0.0768 (0.75288)	-0.513 (0.13495)	-0.437 (0.20333)
Constant	0.748 (0.66367)	3.125** (0.01337)	0.529* (0.07774)	-0.0689 (0.85590)	-0.0818 (0.96068)	4.828*** (0.00049)	0.202 (0.91083)	-0.735 (0.76353)	4.034 (0.22149)	0.534 (0.40156)	1.129 (0.16765)	5.101*** (0.00001)	6.949*** (0.00000)
N	369	367	367	367	185	185	184	185	184	184	184	368	304
R ²	0.080	0.046	0.072	0.067	0.387	0.186	0.158	0.187	0.224	0.204	0.154	0.096	0.137
Panel B: Comparison of the Online and SeqMatch treatments (controls for games orders included)													
Online	0.353 (0.44963)	0.499 (0.15388)	0.0320 (0.70821)	0.0647 (0.52592)	1.971*** (0.00081)	0.871** (0.02722)	0.354 (0.48965)	0.435 (0.53888)	0.501 (0.62290)	0.153 (0.43184)	0.0280 (0.90162)	-0.201 (0.55129)	-0.114 (0.76333)
N	262	261	261	261	132	132	130	132	130	130	130	262	200
R ²	0.010	0.023	0.010	0.011	0.087	0.044	0.029	0.010	0.011	0.007	0.017	0.007	0.007
Panel C: Comparison of the Online and PayPal treatments (controls for games orders included)													
Online	0.566 (0.23520)	0.102 (0.77762)	-0.0320 (0.69658)	0.0465 (0.65696)	1.776*** (0.00338)	0.618 (0.12082)	1.002** (0.03950)	0.641 (0.37804)	2.069** (0.03431)	0.286 (0.11991)	0.325 (0.14835)	-0.475 (0.14064)	-0.589* (0.08408)
N	262	262	262	262	132	132	130	132	130	130	130	262	208
R ²	0.013	0.006	0.007	0.008	0.073	0.027	0.044	0.010	0.047	0.022	0.038	0.011	0.017

Notes: OLS estimates with baseline=Inlab. *p*-values are reported in parenthesis. *, ** and *** denote statistical significance at the 10, 5 and 1% levels. Panel A compares the Inlab treatment to the other three treatments. Demographic controls are all variables from table 4. Beliefs over the experiment controls are all variables from table 5. Game specific timing variables are standardized. . Panels A and B compare the Online treatment to the SeqMatch and PayPal treatments, respectively (constants not reported; regressions control for games ordering effects only). *Public Good Game*: *Contribution* = unconditional contribution to the common project; *Mean conditional contributions* = mean of conditional contributions to the common project; *Slope against low* = slope of the reaction function for average contributions of other group members from 0 to 5; *Slope against high* = slope of the reaction function for average contributions of other group members from 6 to 10. *Dictator game*: *Transfer* = transfer in the Dictator game. *Ultimatum game*: *Transfer* = transfer in the Ultimatum game; *Transfer threshold* = minimum acceptable offer in the Ultimatum game. *Trust game*: *Amount sent* = amount transferred in the Trust game; *Mean amounts returned* = mean of the amounts returned to participant A; *Slope against low* = slope of the reaction function for amounts transferred by participant A from 1 to 5; *Slope against high* = slope of the reaction function for amounts transferred by participant A from 6 to 10. *Holt&Laury lotteries*: *Nb safe choices* = number of times (out of 10) the subject chose the secure option (*i.e.* option A); *Nb safe choices w/o confused* = number of times (out of 10) the subject chose the secure option (*i.e.* option A) excluding the sub-sample of confused subjects, *i.e.* all subjects who either chose the secure option (*i.e.* option A) in the last decision or switched back from option B to option A at least once.

Appendix

Table A.1 Regression analysis (i) excluding all confused subjects in the Holt&Laury task (ii) excluding overbooked subjects in the Online treatment and (iii) controlling for risk aversion in the Holt&Laury task

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	<i>Public Good</i>				<i>Dictator</i>	<i>Ultimatum</i>		<i>Trust</i>				<i>Holt&Laury lotteries</i>	
	Contribution	Mean conditional contributions	Slope against low	Slope against high	Transfer	Transfer	Transfer threshold	Amount sent	Mean amounts returned	Slope against low	Slope against high	Nb safe choices	Nb safe choices <i>w/o confused</i>
Panel A: excluding all confused subjects in the Holt&Laury task													
Online	-0.768 (0.16805)	-0.0364 (0.92244)	-0.0232 (0.82179)	0.107 (0.38338)	1.415*** (0.00214)	0.870** (0.01293)	0.628 (0.16247)	1.333* (0.07163)	2.256*** (0.00976)	0.443*** (0.00667)	0.389* (0.05398)	-0.733*** (0.00634)	.
Constant	-1.018 (0.64158)	1.777 (0.26169)	0.373 (0.39221)	-0.127 (0.80677)	-1.140 (0.61178)	2.869 (0.10804)	0.606 (0.80668)	-2.303 (0.52258)	1.876 (0.67674)	0.342 (0.68392)	1.255 (0.23350)	6.810*** (0.00000)	.
N	207	206	206	206	100	100	107	100	107	107	107	206	.
R ²	0.144	0.098	0.091	0.094	0.271	0.275	0.134	0.232	0.293	0.304	0.281	0.141	.
Panel B: excluding overbooked subjects in the Online treatment													
Online	-0.364 (0.49021)	-0.231 (0.53338)	-0.0356 (0.71206)	0.157 (0.16846)	1.793*** (0.00004)	0.736** (0.04144)	0.603 (0.18607)	1.205* (0.06341)	2.738*** (0.00291)	0.489*** (0.00541)	0.430* (0.05428)	-0.513** (0.04999)	-0.579** (0.04093)
Constant	0.0461 (0.98143)	2.951** (0.04299)	0.674* (0.07561)	0.509 (0.25503)	-1.040 (0.58456)	3.589** (0.03537)	2.032 (0.41186)	-1.355 (0.63671)	7.760* (0.09818)	1.070 (0.23227)	2.289** (0.04856)	4.485*** (0.00095)	6.528*** (0.00002)
N	227	226	226	226	116	116	111	116	111	111	111	226	180
R ²	0.136	0.110	0.097	0.117	0.393	0.191	0.171	0.288	0.263	0.253	0.214	0.139	0.181
Panel C: controlling for risk aversion in the Holt&Laury task													
Online	-0.523 (0.28181)	-0.190 (0.57648)	-0.102 (0.25499)	0.146 (0.18234)	1.799*** (0.00001)	0.659** (0.04615)	0.560 (0.18119)	1.213** (0.04732)	2.499*** (0.00241)	0.457*** (0.00407)	0.405** (0.04233)	.	.
Constant	-0.334 (0.86289)	2.694* (0.05701)	0.622* (0.09402)	0.0577 (0.89857)	-1.388 (0.48260)	3.188* (0.06046)	-0.110 (0.96188)	-0.565 (0.84950)	5.365 (0.20584)	0.637 (0.43634)	2.326** (0.02648)	.	.
N	256	255	255	255	130	130	126	130	126	126	126	.	.
R ²	0.124	0.108	0.102	0.098	0.400	0.181	0.191	0.215	0.272	0.258	0.241	.	.

Notes: Same regressions as in Table 6 – Panel B, but (i) excluding all confused subjects in the Holt & Laury task – *i.e.* all subjects who either chose the secure option (*i.e.* option A) in the last decision or switched back from option B to option A at least once (Panel A) (ii) excluding all overbooked subjects in the online treatment – *i.e.* all subjects who logged-in the online experiment and completed it after the target number of participants had been reached (Panel B) and (iii) controlling for risk aversion levels in the Holt&Laury task.

CHAPTER 2

COOPERATION IN A PEER PRODUCTION ECONOMY: EXPERIMENTAL EVIDENCE FROM WIKIPEDIA

Cooperation in a Peer Production Economy

Experimental Evidence from Wikipedia*

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

May 15, 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, Wikipedia

* 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 outstanding 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 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, Gabriel Smaghue and Marie Claire Villeval. 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@eu.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_2012, accessed February 2013.

⁵ See

http://www.pmlive.com/find_an_article/allarticles/categories/General/2011/june_2011/features/dr_wikipedia_will_see_you_no_w..._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 2014) 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

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).

environments, while altruism may not be. In this process, reciprocity and social image appear as 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 section 6 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

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

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.

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.

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 Design of the study

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

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 dollar invested in the common project by a member of the group yields a return of 0.4 dollar 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.*

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

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

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 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

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*

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.

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.

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

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

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.

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 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).

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

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,

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 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

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).

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).

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.

4 Results

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

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 (2010, 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

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 advertised 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.

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 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

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.

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.

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)).

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.

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

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

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).

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).

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.

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 see 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.

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 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 reassuring 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

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

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 fundraising. *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*.

- Gunthorsdottir, 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., 2014. Social Preferences in the Online Laboratory: A Randomized Experiment. *Forthcoming, Experimental Economics*.
- 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.
- Pancier, 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., 2010. *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

Section 1/4 - Description

In this section, groups of 4 participants (yourself and 3 other participants) are randomly formed.

Remember: The participants who belong to your group in this section are different from those you encounter in the other sections of the study.

At the beginning of this section, each member of the group receives \$10.

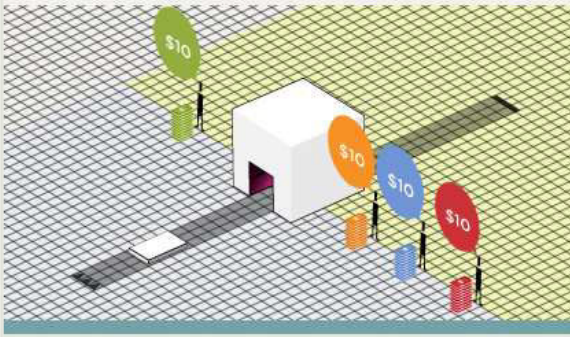
Each member of the group must then decide how many dollars to keep for himself or herself and how many to invest in a common project.

Each dollar invested in the common project by a member of the group yields a return of \$0.40 to each of the 4 group members (including yourself). In other words, the total amount of the contributions to the common project is multiplied by 1.6 before being evenly distributed between the 4 group members.

Your earnings in dollars at the end of this section are given by:


$$10 - (\text{your contribution to the common project}) + 0.4 \times (\text{total contribution to the common project})$$

=> The next screen gives examples...



← Previous Next →

Figure 2. A typical Barnstar



The Teamwork Barnstar
Cas, for being one of those awesome wikipedians who produces great content in a collegial manner, helping out all over, and great dispute resolution. — [Rlevse](#) • [Talk](#) • 13:50, 26 December 2009 (UTC)

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

Section 1/4 - Enter your decision 2/2

This is a decision screen. Once you have made your decision and clicked the "Next" button, you will not be able to go back to this screen again.

* You are now provided with a contribution table that lists each possible average contribution that the other group members could make (all integers between 0 and 10).

For each possible average contribution of the other group members, how much do you want to invest in the common project?

If the other group members make an average contribution of:	\$0	\$1	\$2	\$3	\$4	\$5	\$6	\$7	\$8	\$9	\$10
How much do you want to invest in the common project?	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

[Review description](#) YOU CAN READ THE DESCRIPTION OF THIS SECTION AGAIN AT ANY TIME BY CLICKING HERE

[Previous](#) [Next](#)

Figure 4. The Wikipedia recruitment banner



The banner features a stylized illustration of a person standing in a hallway with a red speech bubble containing a question mark. Logos for the Berkman Center for Internet & Society at Harvard University and SciencesPo are displayed. The main text reads: "Please help advance research with a quick interactive online experiment". A "Learn more now!" button is located in the bottom right corner. At the bottom left, it says "With support from the Wikimedia Research Committee".

BERKMAN CENTER FOR INTERNET & SOCIETY AT HARVARD UNIVERSITY

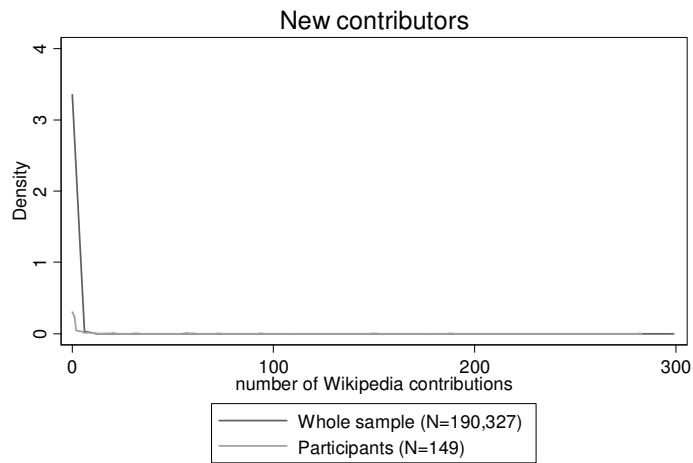
SciencesPo.

Please help advance research with a quick interactive online experiment


With support from the Wikimedia Research Committee

Learn more now!


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



note: the graph discards data above 300 contributions



note: the graph discards data above 100,000 contributions



note: the graph discards data above 400,000 contributions

Notes: Kernel density estimates.

Table 1. Descriptive statistics

	New contributors	Engaged contributors	Administrators	Other
Number of observations (<i>N</i>)	149	566	120	15
DEPENDENT VARIABLE				
Mean – number of Wikipedia contributions	8.64 (3.56) [0; 273]	9719.83 (23519.39) [303; 364157]	41229.24 (86191.33) [2475; 922895]	543.13 (1664.19) [0; 6547]
SOCIAL PREFERENCES MEASURES				
<i>Altruism (N=405)</i>				
(i) Proportion of endowment transferred – Dictator	0.38 (0.30)	0.36 (0.30)	0.42 (0.28)	0.28 (0.26)
(ii) Proportion of endowment transferred – in-group Dictator	0.46 (0.29)	0.45 (0.30)	0.48 (0.28)	0.40 (0.20)
<i>Reciprocity (N=850 & N=445)</i>				
(i) Average proportion of endowment conditionally contributed – Public Goods	0.45 (0.27)	0.54 (0.24)	0.52 (0.24)	0.52 (0.26)
(ii) Average proportion of amount returned – Trust	0.45 (0.25)	0.51 (0.23)	0.46 (0.23)	0.54 (0.22)
<i>Social image (N=850 & N=456)</i>				
(i) Mean size of Wikipedia user page (in bytes)	453.81 (3597.62)	5586.24 (10859.97)	9179.64 (11012.09)	1238.60 (3438.12)
Number of Barnstars receivers	4	340	109	3
Mean - number of Barnstars received	1.5 (0.58)	6.14 (8.57)	16.8 (15.99)	5 (6.93)
(ii) Proportion signaling Barnstars	0.25 (0.50)	0.50 (0.50)	0.70 (0.46)	0.33 (0.58)
DEMOGRAPHIC VARIABLES				
Age	27 (11.81)	34 (14.73)	34 (12.86)	33 (8.84)
Proportion female	0.15 (0.36)	0.09 (0.29)	0.11 (0.31)	0.07 (0.26)
Degree level	3.97 (1.92)	4.55 (1.80)	4.88 (1.64)	4.73 (1.75)
Salary level	3.17 (2.15)	3.80 (2.34)	4.01 (2.25)	3.79 (2.12)
Risk aversion level	6.16 (2.36)	5.66 (2.34)	5.53 (2.38)	4.67 (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
<i>Age</i>		
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%
<i>Gender</i>		
Proportion female	9%	10%
<i>Education level</i>		
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*** (0.00306)	0.0136*** (0.00518)	0.00516 (0.00338)	0.0160*** (0.00541)
female	-0.365** (0.147)	-0.665*** (0.243)	-0.0332 (0.152)	-0.244 (0.249)
degree level	0.0582** (0.0246)	0.0971*** (0.0362)	0.0165 (0.0309)	-0.0391 (0.0446)
salary level	0.00282 (0.0200)	-0.00846 (0.0313)	-0.0133 (0.0225)	-0.0509* (0.0305)
Risk aversion	-0.0325* (0.0169)	-0.0640** (0.0259)	-0.00170 (0.0193)	0.00533 (0.0279)
Constant	8.310*** (0.175)	5.901*** (0.266)	9.420*** (0.189)	10.46*** (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 (0.207)		-0.184 (0.332)		-0.239 (0.224)		0.181 (0.317)	
Altruism (Dictator in-group)		-0.184 (0.207)		-0.177 (0.332)		-0.319 (0.222)		0.199 (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** (0.162)		0.796*** (0.246)		-0.107 (0.187)		-0.631** (0.307)	
Reciprocity (Trust)		0.443* (0.242)		1.136*** (0.392)		0.0424 (0.273)		-0.990** (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:	Whole	Below	Above	Whole	Admins	Admins
number of Wikipedia contributions	sample	median	median	sample		
Social signaler (user page)	1.305*** (0.0845)	1.805*** (0.126)	0.261*** (0.101)		0.354*** (0.135)	
Social signaler (Barnstars)				0.288*** (0.0969)		0.089 (0.151)
nb Barnstars				0.0405*** (0.00287)		0.0319*** (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*** (0.190)	2.021*** (0.281)	2.096*** (0.270)	2.133*** (0.466)	0.678*** (0.230)	0.562 (0.348)		
Social signaler (Barnstars)							0.665*** (0.224)	0.812** (0.335)
Reciprocity (Public Goods) x Social signaler (user page)	-0.0987 (0.198)		0.354 (0.304)		-0.442* (0.228)			
Reciprocity (Public Goods) x non Social signaler (user page)	0.563** (0.250)		0.932*** (0.318)		0.338 (0.321)			
Reciprocity (Trust) x Social signaler (user page)		-0.294 (0.309)		0.632 (0.586)		-0.0176 (0.316)		
Reciprocity (Trust) x non Social signaler (user page)		1.067*** (0.376)		1.173** (0.501)		0.451 (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*** (0.00288)	0.0371*** (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 (0.180)	-0.0265 (0.272)	-0.0393 (0.187)	-0.730** (0.309)	-1.004** (0.463)	-0.725* (0.419)	-0.626 (0.460)	-3.703* (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"."

Appendix

This section checks the robustness of our results with alternative measures of social preferences. Building upon Fischbacher et al.'s approach (2001), we classify subjects into four different groups depending on whether they reveal a preference for reciprocity or altruism in the conditional Public Goods game. We compute (i) the slope of subjects' reaction functions to the possible average contributions of the other group members ("reciprocity") and (ii) the average proportion of the endowment that is conditionally contributed across all 11 conditional contributions decisions ("mean contribution"). We then classify subjects according to the following rule:

- Free riders: Reciprocity < 1 & Mean contribution < 0.2
- Weak reciprocators: reciprocity < 1 & $0.2 < \text{mean contribution} < 0.8$
- Reciprocators: reciprocity > 1
- Altruists: reciprocity < 1 & Mean contribution > 0.8

This classification distinguishes between "weak" and "non weak" reciprocators in order to remain consistent with the typology initially proposed by Fischbacher et al. (2001). The distinction between those two types of reciprocators can be important to understand the sustainability of contributions to Wikipedia. As shown in the lab, groups constituted of reciprocators usually succeed in sustaining their contribution levels in repeated Public Goods experiments. On the other hand, the presence of free-riders and weak reciprocators in a group usually triggers the progressive decline of cooperation that is typical of lab experiments. An additional value-added of this classification is to verify that our results are robust to including all three preferences in the same regression.

Table 10. Alternative measures of social preferences and robustness check for our main results

	(1)	(2)	(3)	(4)
Dependent variable: number of Wikipedia contributions	Whole sample	Below median	Above median	Admins
Free rider (Public Goods) x Social signaler (user page)	1.538*** (0.240)	2.278*** (0.382)	0.651** (0.256)	0.385 (0.372)
Weak reciprocator (Public Goods) x Social signaler (user page)	1.519*** (0.198)	2.681*** (0.307)	0.286 (0.220)	0.0473 (0.334)
Weak reciprocator (Public Goods) x Non social signaler (user page)	-0.0683 (0.214)	0.513 (0.313)	-0.00447 (0.266)	-0.437 (0.334)
Reciprocator (Public Goods) x Social signaler (user page)	1.541*** (0.197)	2.438*** (0.316)	0.288 (0.218)	-0.115 (0.330)
Reciprocator (Public Goods) x Non social signaler (user page)	0.421** (0.204)	1.067*** (0.305)	-0.00849 (0.240)	-0.319 (0.361)
Altruist (Public Goods) x Social signaler (user page)	1.459*** (0.222)	2.644*** (0.345)	0.214 (0.246)	-0.352 (0.454)
Altruist (Public Goods) x non Social signaler (user page)	0.487** (0.241)	1.019*** (0.351)	0.301 (0.282)	-0.530 (0.425)
Control variables	Yes	Yes	Yes	Yes
N	649	325	324	113
Pseudo R ²	0.0269	0.0521	0.00242	0.00889

Notes: Negative binomial estimates. Standard errors are reported in parenthesis. Constants not reported. *, ** and *** denote statistical significance at the 10, 5 and 1% levels.

CHAPTER 3

SOCIAL MOTIVES AND THE ORGANIZATION OF PRODUCTION: EXPERIMENTAL EVIDENCE FROM OPEN SOURCE SOFTWARE

Social Motives and the Organization of Production: Experimental Evidence from Open Source Software*

Yann Algan[†] Yochai Benkler[‡] Emeric Henry[§] Jérôme Hergueux[¶]

May 15, 2014

Abstract

Most organizations rely not only on monetary incentives, but also on a variety of other motives to elicit effort and encourage innovation. We rely on the community of open source software (OSS) developers to study how heterogeneous motivations affect individual contributions and endogenous team formation. Our primary focus is on the role of social motivations, which we elicit experimentally on a stratified sample of 1,194 OSS developers. We show that social motives predict individual contributions as strongly as extrinsic ones, but have a different pattern of association with individual contributions. Socially motivated developers have a lower extensive margin of contributions (i.e. they seek affiliation with less projects) but a significantly higher intensive margin, with an overall positive association with contributions. We also report evidence that cooperative developers tend to match assortatively in development teams, while non cooperative ones seek to join teams that are comprised of different social types than their own. Cooperative project administrators act as the gatekeepers of their development teams, and tend to screen developers based on their social type. Generally speaking, this paper is a first attempt at showing how experimental games can help explain productive activities and industrial organization.

JEL Classification: H41, C93, D23, L23

Keywords: Field Experiment, Public Goods, Social Preferences, Peer Production, Open Source Software, Labor Economics

*We gratefully acknowledge financial support from the European Research Council (ERC Starting Grant) and logistical support from the Sciences Po médialab. We are greatly indebted to Sourceforge.net for endorsing our study and the Sourceforge Research Data Archive (SRDA) at the University of Notre Dame for granting us privileged access to their data. We are grateful to Anne l'Hôte, Romain Guillebert and Matt Van Antwerp for outstanding research assistance. This paper has benefited from invaluable conversations with Shane Greenstein, David Laibson, Gabriel Smagghue, and Jack Willis. We thank the Berkman Cooperation group and numerous seminar participants for helpful comments. Our last thank goes to the 1,194 open source software developers who took part in this study.

[†]Sciences Po, Department of Economics and CEPR. e-mail: yann.algan@sciencespo.fr

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

[§]Sciences Po, Department of Economics and CEPR. e-mail: emeric.henry@sciencespo.fr

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

1 Introduction

Many organizations strongly rely on multiple motivations beyond monetary incentives to elicit effort from their members. This is traditionally the case for public organizations and private non profits aimed at providing goods that are deemed socially beneficial (e.g. fundamental research, public health, education). It is also increasingly the case in for profit firms, particularly in innovative sectors in which it is difficult to contract upon workers' level of effort. Managers of such organizations often strive to forgo monitoring costs and agency problems by designing intrinsically motivating work environments in which workers enjoy a great deal of autonomy. Such organizational strategies can lead to large welfare gains (Bartling, Fehr, and Schmidt, 2012; Bénabou and Tirole, 2013). This increasing reliance on workers' heterogeneous motivations raises numerous questions, however. What motivations best account for individual effort and production? How do those motivations interact to determine individual output? Do certain motivations affect the nature of individual contributions differently? How should workers with heterogenous motivations be grouped to produce more efficiently?

This paper focuses the analysis on the role of social motivations and relies on the community of open source software (OSS) developers to answer the above questions. We stratify all the active software projects registered at Sourceforge.net (by far the largest online OSS development platform at the time of this study) according to their size and license type, and recruit 1,194 developers from those diverse projects to participate in our study. Rather than relying on self-reports, we elicit the social motivations of developers with an online experiment and complement our design with standard survey questions aimed at measuring motives that are not easily captured experimentally. We then relate developers' heterogenous motivations to detailed real-world records of individual contributions to OSS, which we extract separately from the Sourceforge website. Our combination of experimental and field data also allows us to report on how developers seek to match by cooperative type at the team level, a question that the previous literature had only been able to address in a laboratory context.

Open source software provides a unique environment to perform such tests. First, the individual contributions of developers to the final output are observable. Second, those contributions span different types of activities, which allows to study how developers self-select into performing different roles based on their motives. Third, many software development teams with different organizational structures simultaneously coexist and enjoy varying degrees of success.

Not only is OSS a great environment to answer these essential questions, it is also of interest in and of itself. Over the past decade, OSS has moved from an organizational curiosity to becoming the flagship of peer production, a large-scale, collaborative and primarily voluntary based model of innovation and production (Maurer and Scotchmer, 2006; Benkler, 2013). Peer production therefore represents an emerging form of organization in which individuals voluntarily self-assign work and successfully coordinate towards the provision of global public goods, in the absence of price signals and without any pre-specified design rule or formal leadership.¹ OSS

¹The open source development model was actually a major source of inspiration for firms such as Google and Mi-

currently involves an estimated 800,000 developers around the world (Crowston, Wei, Howison, and Wiggins (2012)). It is also a major source of innovation. OSS is responsible for most of the basic utilities on which the Internet runs (e.g. the Domain Name System management software BIND, Sendmail, the Apache web server), popular programming languages (e.g. Python, Perl) and programming environments (e.g. Eclipse). OSS also successfully competes with many of its firm-based counterparts in the realm of end-user applications, and even enterprise systems. Linux, for instance, has been adopted by a number of large commercial firms such as IBM, Apple and Sun for its relative reliability, resilience to virus attacks, and bug correcting speed.²

One of the enduring questions is, as Lerner and Tirole phrase it: "why would thousands of top-notch programmers contribute freely to the provision of a public good" Lerner and Tirole (2002). Some models emphasize social motives as the key drivers (Bitzer, Schrettl, and Schröder (2007)), while others propose monetary compensation and reputation building, grounded in standard economic theory (see Lerner and Tirole (2002)). However the empirical literature has not succeeded in providing a clear picture of developers' motivations. This is particularly important since the nature of developers' motivations can largely determine the circumstances under which the peer production model dominates the proprietary model³ and whether the emergence of peer production is welfare enhancing.⁴

Microsoft when they initially designed working environments that emphasize individual autonomy and heavily leverage intrinsic motivations to work and innovate (e.g., reciprocity, altruism, peer recognition) Kogut and Metiu (2001). Indeed, firms have recognized the potential of peer production as an organizational model early on, either by directly supporting OSS projects (i.e. allocating a fraction of their labor force to its development) or trying to leverage peer production communities to develop their own products in an open source mode. (See Hars and Ou (2001), Lerner and Tirole (2005a), Maurer and Scotchmer (2006) and Jullien and Zimmermann (2009) for an account of how firms benefit from pursuing such open source strategies.) At the turn of the century, IBM notably announced a strategic decision to invest over 1 billion dollars to support the development of OSS, while a third of the world's 25 largest software companies were already engaged in significant open source activities (Ghosh, Glott, Krieger, and Robles (2002)). More 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.

²Indeed, Walli, Gynn, and Rotz (2005) report that 87% of US businesses now rely on OSS for some of their daily activities. Ghosh (2007) estimates the cost of recreating the existing open source code at 12 billion euros, while Greenstein and Nagle (2014) estimate the value of the Apache web server alone to range between 2 and 12 billion dollars.

³Kuan (2001), Gaudeul (2004) and Bessen (2005), for instance, develop models in which OSS is better able to serve sophisticated users, while proprietary software is better able to serve unsophisticated users. Indeed, OSS allows sophisticated users to tailor the program to their specific needs, thus creating a rich and freely accessible common pool of customized features. At the same time, OSS developers lack the proper incentives to develop easy to use interfaces and provide services that are appealing to unsophisticated users (e.g., documentation). On the other hand, proprietary software manufacturers cannot anticipate all possible manifestations of consumer demand. They therefore tend to focus on developing easy to use interfaces in order to reach the broader pool of unsophisticated users and increase the market penetration of their products. Such conclusions, however, are only valid insofar as OSS developers' motivations to contribute are mainly intrinsic - non social (e.g. "own-use"), as opposed to intrinsic - social (e.g. "altruism").

⁴Schmidt and Schnitzer (2002) and Saint-Paul (2003) both develop models in which open source software becomes welfare reducing in the long-run if developers' intrinsic motivations to contribute are significantly weaker than their extrinsic motivations. From a static point of view, OSS increases social welfare. However, in a dynamic view, the emergence of OSS drastically reduces the profitability of the software industry in the long-run, thus reducing OSS

There are three main limitations in the existing literature on developers' motivations: (i) the overwhelming majority of the existing empirical research relies on self-reports to elicit motivations and does not try to relate those stated motives to objective measures of contributions to OSS, (ii) the existing literature typically relies on small sample surveys of developers, with a strong sampling bias towards a handful of very big and successful projects, which prevents the literature from getting a more representative picture of the motivations of the overall population of OSS developers and (iii) the empirical research on developers' motivations has typically not been grounded in economic theory.⁵ This paper seeks to concurrently address those three limitations in the current OSS literature.

We elicit a variety of potential motivations for individual developers to participate in OSS development. Among those motives, our main focus is on the role of the private social motivations that we elicit experimentally. Those motivations are the ones that the theoretical literature on the private provision of public goods has put forward to account for individual's often observed willingness to sustain cooperation in public goods like environments, be it in the laboratory or in the field: (i) altruistic motives, either in the form of "pure altruism" or "warm-glow" (Andreoni, 1989, 1990; Anderson, Goeree, and Holt, 1998) and (ii) reciprocity motives (Rabin, 1993; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006).

On top of those experimentally elicited motives, we also consider the potential role of two additional intrinsic – but non social – motives to contribute, which are put forward by the literature in labor economics: (i) own use (i.e. some developers may have very specific needs which are not met by any market-based offer, see Lerner and Tirole (2002)) and (ii) the pleasure derived from learning and problem solving (i.e. developers might have a positive marginal utility of effort over some range, see Kreps (1997) and Glazer (2004)). Those potential intrinsic motivations go along with more traditional extrinsic motives: (iii) reputation motives, e.g. in the form of future monetary rewards through labor market signaling (Spence (1974)) and (iv) immediate monetary payoffs. Indeed, firms' early involvement in supporting and developing OSS explains that around 50% of OSS developers consistently declare receiving direct or indirect monetary payments to contribute across existing survey studies (see Ghosh, Glott, Krieger, and Robles, 2002; Hertel, Niedner, and Herrmann, 2003; Lakhani and Wolf, 2005). Finally, one additional intrinsic motive deserves some scrutiny even though it has not been explicitly taken into account by economic theory, for it is the first motivation to contribute to OSS put forward by the open source movement's historical leaders (e.g. Richard Stallman): the role of ideology and the belief that "information wants to be free" and, consequently, that software should not be a proprietary product.⁶ While we mainly think of those self-reported motivations as control variables, testing for their association with developers' real world contribution patterns is interesting in its own right. Indeed, the existing

developers' extrinsic incentives to contribute. Such a conclusion, again, is only valid insofar as the assumptions made about the nature of developers' motivations to contribute are themselves valid.

⁵Indeed, most of the existing empirical literature relies on theoretical frameworks imported from management studies and social psychology.

⁶See the GNU manifesto: <https://www.gnu.org/gnu/manifesto.html> and "why software should not have owners": <https://www.gnu.org/philosophy/why-free.html>, accessed March 2014.

literature has been relying heavily on such survey items to learn about developers' motivations for over a decade, without actually having tried to correlate them with field data.⁷

We find that developers rank intrinsic motivations the highest in the survey. However, those motivations, if anything, often weakly correlate with developers' real-world contributions records. Conversely developers rank the motivations that have an extrinsic component (i.e. reputation concerns and monetary compensation) the lowest. Those motivations are nonetheless the self-reported ones that best correlate with their individual contributions. Among our population of developers – and controlling for all other factors – the social motives that we elicit experimentally are the only ones that are as strongly associated with participation as extrinsic ones. They have a different pattern of association with individual contributions, however. Consistent with a labor market signaling hypothesis, developers who want to get a reputation have a significantly higher extensive margin of contributions (i.e. they seek to get formal affiliation with more projects, which is a signal of quality that can be easily observed by external firms). By contrast, developers who are socially motivated have a significantly lower extensive margin, but a very high intensive margin of contributions (with a net positive effect on contributions). Interestingly, we report evidence that cooperative developers tend to match assortatively in development teams. By contrast, self-interested developers seek to join teams that are comprised of different social types than their own. This assortative matching effect appears as mainly driven by the project administrators, who effectively act as the gatekeepers of their development teams and seek to coopt developers of their own social type.

This paper is related to two main literatures. The first one looks at the intrinsic (Benkler (2013)) versus extrinsic (Lerner and Tirole (2002)) motivations that underpin peer production. The second literature looks more generally at social preferences within organizations. So far this literature is mostly theoretical. Seminal articles studied the heterogeneity of preferences and workers' endogenous sorting behavior by missions (Besley and Ghatak (2005)) or motivations (Kosfeld and von Siemens, 2011). They predict the absence of pooling of types within organizations, as free-riders crowd out cooperative types. On the empirical side, the evidence comes from lab experiments (Falk and Kosfeld, 2006; Bartling, Fehr, and Herz, 2013). A few papers look at real-world organizations, but focus on particular proxies for social motives, such as workers' productivity under various incentive schemes (see Bandiera, Barankay, and Rasul (2005)). Our paper distinguishes itself from those literatures by being the first to elicit agents' social motives by bringing experimental economics within productive organizations. On the methodology side, this paper illustrates the 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 and diverse samples and achieve high internal validity with online experiments (Hergueux and Jacquemet (2014)), the Internet also provides a wealth of externally valid observational data on individuals' field behavior (Lazer, Pentland, Adamic, Aral, Barabasi, Brewer, Christakis, Contractor, Fowler, Gutmann, et al. (2009)).

⁷Roberts, Hann, and Slaughter (2006) is the only exception that we know of.

The rest of the paper proceeds as follows. Section 2 documents the design and implementation of the study, together with our dependent variables and identification strategy. We report our developer level results on the effect of heterogeneous motivations to work in section 3, and study the team level sorting behavior of individuals in section 4. Section 5 provides a discussion of our results and concludes.

2 Design of the study

In this section, we first describe our strategy for measuring motives among our subjects. We then describe our experimental procedures before reporting on the practical implementation of the experiment. We end this section by presenting our main dependent variables and identification strategy.

2.1 Measuring motivations

Since OSS is a privately provided public good (Lerner and Tirole, 2002; Bitzer, Schrettl, and Schröder, 2007), economic theory suggests that social motivations could play a significant role in incentivizing individual contributions. The literature in labor economics and industrial organization, however, has identified other intrinsic (and not necessarily social) motives for contributing, together with traditional extrinsic ones. We first rely on an experimental approach to elicit social motives within our sample of OSS developers. In a second step, we build upon the previous literature's survey approach to elicit other intrinsic and extrinsic motives.

2.1.1 Experimental measures

We elicit social motives among our subjects using experimental data from four mostly standard decision problems taken from the literature on social preferences (see, e.g., Fehr and Camerer (2004)).

Public Goods Game. Following Fischbacher, Gächter, and Fehr (2001), we use a modified version the Public Goods game to elicit subjects' altruism and reciprocity motives. 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 dollar invested in the common project by a member of the group yields a return of 0.4 dollar to each group member.⁸ 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

⁸Each subject thus faces the following payoff function:

$$\pi_i = 10 - contrib_i + 0.4 \sum_{j=1}^4 contrib_j \quad (1)$$

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.

Building upon Fischbacher, Gächter, and Fehr (2001)'s approach, we classify subjects into four different groups depending on whether they reveal a preference for reciprocity or altruism in the conditional Public Goods game. We compute (i) the slope of subjects' reaction functions to the possible average contributions of the other group members ("reciprocity") and (ii) the average proportion of the endowment that is conditionally contributed across all 11 conditional contributions decisions ("mean contribution"). We then classify subjects according to the following rule:

1. *Free riders*: reciprocity < 1 & mean contribution ≤ 0.2
2. *Weak reciprocators*: reciprocity < 1 & $0.2 < \text{mean contribution} < 0.8$
3. *Reciprocators*: reciprocity ≥ 1
4. *Altruists*: reciprocity < 1 & mean contribution ≥ 0.8

This classification distinguishes between "weak" and "non weak" reciprocators in order to remain consistent with the typology initially proposed by Fischbacher, Gächter, and Fehr (2001).⁹ Furthermore, distinguishing between those two types of reciprocators can also be important at a theoretical level: groups constituted of reciprocators usually succeed in sustaining their contribution levels in a repeated Public Goods experiment. On the other hand, the presence of free-riders and weak reciprocators in a group usually triggers the progressive decline of cooperation that is typical of lab experiments (Chaudhuri (2011)).

Dictator game. The Dictator game is certainly experimental economics' workhorse for studying altruistic motives. As a result, we use a standard Dictator game on top of our main above approach in order to elicit this preference among our subjects.¹⁰ Each participant A is matched with a participant B to play as a dictator. The dictator receives a \$10 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 developers are incentivized to contribute to OSS out of altruism directed towards their fellow contributors, we provide an alternative measure for this motive by conducting two additional Dictator games in which we induce some in-group bias. We do this by telling subjects that they are now matched with another subject who (i) "participates in open source projects" and (ii) "par-

⁹In their setting, our "weak reciprocators" would be labeled "hump shaped" contributors.

¹⁰Note 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.

participates in open source projects with a GPL license".¹¹ *We take the proportion of the endowment transferred by the dictator in both directed decisions as alternative measures of subjects' altruistic motive.*

Trust game. The Trust game initially proposed by Berg, Dickhaut, and McCabe (1995) has been extensively used in the literature to elicit reciprocity preferences. As a result, we also use a standard Trust game in order to elicit this preference 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.*

Ultimatum Bargaining game. We use this game in order to elicit subjects' fairness preference. Fairness considerations are notably likely to matter for the choice of the project license (see, e.g., Lerner and Tirole (2005b)). Each participant A is matched with a participant B. Participant A is the proposer and must decide on how much of an initial endowment of \$10 is transferred to participant B - the responder. The responder is simultaneously asked for the threshold level of transfer below which the offer will be refused. The earnings of each player in this game are computed according to the proposal if participant A's transfer is higher or equal to the threshold. Otherwise, both players' earnings are set equal to 0. *We take the proportion of the endowment demanded by the responder in the Ultimatum Bargaining game a measure of his fairness motive.*

2.1.2 Survey measures

The literature on open source software has a long tradition of relying on survey questions to elicit developers' motivations to contribute to OSS and compare their prominence. Since Ghosh, Glott, Krieger, and Robles (2002)'s seminal survey of OSS developers and the subsequent follow-ups by David, Waterman, and Arora (2003) and Shimizu, Iio, and Hiyane (2004), the survey items used to elicit those motives have remained much the same. We thus include some of those standard items in our post-experimental questionnaire in order to (i) measure the intrinsic and extrinsic motivations that are difficult to elicit experimentally (ii) test for the predictive power of those survey measures of motives on developers' real world contribution records. Specifically, we ask subjects to state their level of agreement with the following reasons for contributing to OSS:

1. Because I think that software should not be a proprietary product (motive = ideology)
2. Because I like to learn and develop new skills (motive = like to learn)
3. Because I need to solve a problem that could not be solved by proprietary software (motive = own use)

¹¹The General Public License (GNU) is a famous and prominent copyleft license. It guarantees that end users retain the freedoms to use, study, share, and modify any derivative work.

4. Because I want to get a reputation in the OSS developers scene (motive = establish reputation)
5. To make money (motive = earn money)

Motivations 1 to 3 above are intrinsic - but non social - motivations to contribute. The "ideology" motive was the first to be put forward by the OSS movement's historical leaders, such as Richard Stallman. The "like to learn" and "own use" motives have been both formalized in the labor economics and industrial organization literature. Motivation 4 measures subjects' social image motive, jointly in the form of peer recognition and of future monetary rewards through labor market signaling. It can therefore be considered intrinsic or extrinsic, depending on whether subjects seek to establish their reputation in the OSS community out of an intrinsic taste for peer recognition or in order to enhance their future labor market outcomes. Indeed, both subcomponents of the social image motive lead to similar theoretical predictions in terms of effects, and so it makes sense to use this broad formulation to consider their joint impact on OSS contributions (see, e.g., Bénabou and Tirole (2006)). Finally, motivation 5, "earn money", is the paradigmatic extrinsic motivation, i.e. direct monetary payoffs. At the end of the survey, we also ask subjects whether they actually are getting paid to contribute to OSS, be it directly or indirectly (as opposed to merely stating that making money is an important reason why they contribute to OSS).

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, Falk, Huffman, Sunde, Schupp, and Wagner (2011).

2.2 Experimental procedures

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 OSS oriented. Importantly, we made it very clear on the introductory screen that subjects would interact with a diverse pool of Internet users.¹² 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.¹³ We only require a valid e-mail address to process the payment. To strengthen the cred-

¹²The OSS subjects were matched with a traditional pool of laboratory subjects who had previously participated in a similar online experiment.

¹³Such 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.

ibility 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 OSS developers can be very hostile to monetary rewards.¹⁴ In order to ensure that the experiment is equally incentive compatible for all subjects, we allow them to donate their final earnings to the International Committee of the Red Cross - a renowned and general purpose charitable organization - 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 seven decisions, followed by the survey, are made successively following a given sequence of screens. The unconditional and conditional Public Goods games are the most important to our design but also 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 games with induced in-group bias to generate spillover effects on the other decisions, we always present those two decision problems at the very end of the experiment. We always present the standard Dictator game first, the Dictator game with OSS in-group bias second and the Dictator game with OSS-GPL in-group bias third. Finally, in order to alleviate anchoring effects, we sequentially vary the order in which the Trust game and the Ultimatum Bargaining game are presented to subject according to their login order. This led us to implement the two following games ordering:

- Order 1: *Public Goods - Ultimatum - Trust - Dictators*
- Order 2: *Public Goods - Trust - Ultimatum - Dictators*

All decisions made by our subjects are anonymous. This is because all the preferences that we elicit experimentally are private social preferences, meaning that they do not depend on the visibility of one's actions to be at work. 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).¹⁵ Second, the instructions screens are

¹⁴In their survey based study of OSS developers motivations, Haruvy, Wu, and Chakravarty (2005) notably report that "quite a few respondents sent e-mails expressing indignation at survey items which suggested monetary considerations could possibly motivate their contributions to Open Source projects."

¹⁵The 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.

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 2 for the example of the conditional Public Goods game decision screen).

2.3 Implementation of the experiment

The participants to this study are all OSS developers registered with Sourceforge.net. With 221,802 projects registered in 2010, Sourceforge was by far the largest centralized online platform for developers to control and manage free and open source software development at the time. End users can download the various projects directly from the website. Sourceforge provides developers with a comprehensive set of free development tools. Each project notably has a bug tracker system that allows developers to report, document and track the bugs that affect the software and a message forum that can be used by end-users in order to ask questions and request technical support from the community of developers.

There is great heterogeneity within Sourceforge in both the number of contributors that OSS projects attract (i.e. their "size") and the restrictiveness of the licenses that they choose to adopt. As noted by Belenzon and Schankerman (2008), two main features define the restrictiveness of a license: (i) the extent to which the code and any of its modifications can be subsequently embodied in commercial software and (ii) whether modifications to the code have to remain open source (i.e. free to use, study, share, and modify by anyone). In order to get as representative a sample of the population of developers that distributes itself among the universe of peer production efforts, we follow a stratification strategy by project size and license type in order to define our sample of eligible subjects. We select from Sourceforge all the projects that were active in 2010, as defined by having either a bug closed or a last feature added in 2010. This yields a sample of 1,577 active projects. After excluding the projects for which the SVN logs were inaccessible - i.e. the logs to the software revision control system which provides detailed information on code contributions at the developer level - we are left with a sample of 1,242 active projects.

From the 8,858 developers affiliated with those active projects, we identify those who had some development activity in 2010. We then order projects according to their number of active developers, and rely on Belenzon and Schankerman (2008)'s classification of the 44 existing OSS license types to label their licensing terms as highly, moderately or weakly restrictive. Since there are only

83 projects with more than 7 active contributors, we select all of those project irrespective of their license terms. For all the projects with 6 or less active contributors, we choose to construct a sample including an equal number of highly, moderately and weakly restrictive licenses. For instance, out of the 365 projects that have only one active developer, 239 projects feature highly restrictive licenses, 57 projects feature moderately restrictive licenses and 69 feature weakly restrictive licenses. We thus retain the 57 projects with moderately restrictive licenses and then randomly select 57 project from the pool of projects with both highly and weakly restrictive licenses. We end up with a balanced sample of 322 active projects, both in terms of size and license restrictiveness.

All of the 1,019 developers who were active in 2010 on the above 322 projects were eligible to participate in the study. Both for representativity and identification purposes, we also wanted to capture some non contributing developers in our sample of subjects, however. Indeed, the number of contributions made to OSS projects by Sourceforge users follows a strong power law distribution, and many developers make very few contributions, if any.¹⁶ As a result, we also randomly select 3 non contributing developers per project to be eligible to participate in the study. We therefore end up with a stratified sample of 2,534 Sourceforge developers eligible to participate in the study. Table 1 summarizes our approach for constituting this pool of eligible subjects.

With the support of the Sourceforge platform, we collected the e-mail addresses of all eligible developers and sent them individual invitations to participate in the study. Upon clicking on a link included in the invitation message, eligible users were able to log into the system with their Sourceforge username, which allowed us to uniquely identify them and subsequently collect their entire contribution history to OSS. Subjects were then redirected to the welcome screen of the experimental economics platform. Our system sequentially allocated subjects according to their login order, first to the role of participant A or participant B, and second to one of the two possible games ordering. 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 in May 2011 and remained open for 10 complete days. It recruited 1,194 subjects, which establishes the participation rate at 47%.

2.4 Dependent variables and identification strategy

Sourceforge remains the largest web-based repository of OSS development activity in the world to date. It is possible to extract from the platform the complete record of developers real world contributions to OSS. We use this extensive data source to study how developers' heterogeneous motivations determine 3 distinct types of software development activities:

¹⁶For instance, Ghosh, Glott, Krieger, and Robles (2002) report that 10% of Sourceforge developers tend to create 74% of the code, while Mockus, Fielding, and Herbsleb (2002) report that 15 core developers make 15% of all contributions to the Apache project. This feature extends beyond code contributions. Lakhani and Wolf (2005) notably report that 2% of open source tech support providers supply 50% of all answers to end users. Skewness of participation actually characterizes many technology mediated peer production systems and is not unique to OSS. It is also a structural feature of individual contributions to Wikipedia and participation in online message boards.

1. Contributions of code to projects with which the developer is formally affiliated.

The standard way for a developer to contribute code to a project is to request from one of the administrators of the project that he gives him write access to its repository, meaning that he will be able to implement any modification he wants directly to the source code of the project. Modifications made by a developer to the source code of a project are called "commits". A commit is typically a set of changes to the source code that affects several software files at the same time and makes logical sense. They are meant to implement new features or solve existing bugs. A developer who wants to work on a given project will typically start by e-mailing a few code contributions directly to the project administrators. If those contributions are deemed valuable by the project administrators, they will grant the developer "committer status". He will then be automatically referenced as a developer on the list of members of the project featured in its associated Sourceforge web page. Project administrators can therefore be usefully thought of as the gatekeepers of their project, in the sense that they get to decide who becomes a member of the development team. They can also promote a project developer to the status of administrator, in which case he will be able to make changes to the settings of the project (such as its licensing terms) and, in turn, grant committer status to other developers.

We compute two measures of contribution of code to projects with which our subjects are formally affiliated:

- The total number of commits of code made by each subject (variable = nb commits).
- The total number of files added, deleted, modified or replaced by each subject through his commits of code (variable = nb files).

In order to do so, we download from Sourceforge the SVN source code repositories of all the registered projects as of the launch of the experiment. We then analyze the 13,000,000 lines of commit reports together with the 45,000,000 lines of files change reports coming from the 90,000 non empty project logs in order to precisely identify the code contributions of our subjects.

2. Bug resolution attempts on projects with which the developer is not formally affiliated.

Debugging software is an activity that lies at the core of OSS development. As opposed to creating new code, Bessen (2005) reports that testing, debugging and maintenance accounts for 82% of overall software development costs, while Kogut and Metiu (2001) estimate this number to be between 50 and 80%. Since most OSS development teams are rather small, however, they may sometimes not have the human capital necessary to easily solve a particular bug within their team. Some knowledgeable developer from another team could then decide to help and provide a solution. As he does not have write access to the repository of the project, such a developer cannot simply commit a piece of code that would solve the problem. He would therefore need to post his bug resolution proposal (called a "patch") directly on the corresponding thread of the bug tracker.¹⁷ He will also need to provide some detailed documentation together with his patch

¹⁷A developer interested in getting write access to the repository and become formally affiliated with the project

in order to let the developers of the project know how to implement his solution. This is a costly process to go through for a developer (even more so that OSS developers tend not to enjoy writing documentation for their code (Henkel and Tins (2004)), but it can also save significant resources on the part of the development team that is faced with a problem for which it has no handy solution. In order to get a measure of subjects' propensity to help other development teams solve bugs on their own projects, we count the number of messages that they post directly on the bug tracker of a project that have an attachment with a ".patch" or a ".diff" extension. (Those extensions typically denote an attempt at solving a bug.) We therefore compute the following variable:

- The total number of bug resolution proposals posted directly by subjects on the bug trackers of projects (variable = nb patches).

3. End user support activities.

Beyond code contributions, a complete software development and diffusion process requires the performance of "mundane" but important tasks, such as providing end users with technical support and answering their questions about available and non available features. Such end user supporting tasks are typically conducted on the message forum of the project. They are generally considered cumbersome by the majority of developers¹⁸ and can be very time consuming. In their survey study of the famous Apache project, Lakhani and Von Hippel (2003) notably report that the "total annual time spent by information providers [...] at the Apache help forum averages over 100 hours", with the result being that "web server users rank Apache technical support overall as somewhat better than that of its major commercial rivals in the server software field." We compute the following variable in order to get a measure of subject's willingness to engage in end user support:

- The total number of forum messages posted by subjects which follow-up on some initial post (variable = nb messages).

Our 4 main dependent variables above all follow 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 those dependent variables, as this would induce substantial bias in our estimates (Silva and Tenreyro (2006)). As a more cautious approach, we use the negative binomial pseudo-maximum likelihood estimator, which is not affected by this problem.¹⁹ 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.²⁰

is more likely to address his contribution directly to the project administrators who can grant him such rights. By addressing it the entire team of developers he takes the risk that his contribution goes unnoticed by the administrators.

¹⁸As an illustration, Lakhani and Von Hippel (2003) report that the Apache Development Group (a group a core volunteer developers who guide the development and extension of the Apache project) issued a statement in its 1999 FAQ that as developers, they did not want to be "swamped by a flood of trivial questions that can be resolved elsewhere."

¹⁹See chapter 9 in Wooldridge (2010). Log-linearizing our dependent variables to run OLS regressions yields qualitatively similar results (tables available from the authors upon request).

²⁰An alternative estimator that has similar properties is the Poisson pseudo-maximum likelihood estimator. One lim-

3 Developers' motivations to contribute to open source

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 socio-demographic characteristics and patterns of contribution to OSS. Controlling for those socio-demographic factors, we then analyze the association between developers' self-reported motives and their real-world contribution records and test for the predictive power of the pro-social motivations that we elicit experimentally in the context of a conditional Public Goods game. We end this section by analyzing the determinants of developers' intensive and extensive margins of contributions.

3.1 Descriptive statistics and analysis

Table 2 provides some descriptive statistics on our stratified sample of 1,194 developers on (i) their overall number of OSS contributions by type of contribution, (ii) our experimental measures of social motives and (iii) our demographic variables. We can see that the population of OSS developers is young on average (32 years old) and overwhelmingly male (only 3% of developers are female). The average developer has a 4 years college degree (BA, BS), with 17,5% of the population of developers having less than a 2 years college degree and almost half of the population (i.e. 49%) having a Master's or a PhD degree. The average developer earns between \$2,000 and \$4,000 per month, with 32% of the population earning less than \$2,000 and 20% earning more than \$7,500. Overall, those socio-economic characteristics are very consistent across survey studies of OSS developers (see, e.g., David and Shapiro (2008) for a brief review).

Figure 4 features the distribution of cooperative types within our population of developers, which we derive from their revealed preferences in the conditional Public Goods game. We classify subjects into four types: free riders, weak reciprocators, reciprocators and altruists, and vary the "mean contribution" benchmark used to do so from 0 to 0.4. A benchmark of 0 typically means that a subject is classified as a free rider if the average proportion of his endowment that he conditionally contributed across all 11 conditional contributions decisions is equal to 0% (i.e. the subject never makes a positive contribution, irrespective of the average contribution of the other members of the group). Consequently, a subject is classified as an altruist (i.e. an unconditional cooperator) if this average proportion is equal to 100% (i.e. the subject always contributes all of his endowment, irrespective of the average contribution of the other members of the group). Using a benchmark of, say, 0.1 would mean replacing the above thresholds by 10 and 90%, respectively, thus allowing subjects to make "errors" in their contribution decisions by having a looser definition of types.

The main takeaway from this figure is that, depending on the benchmark used, between 7 and 21% of subjects are classified as "altruists", meaning that they decide to *unconditionally* contribute

itation 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.

a very high proportion of their endowment. Such a pattern of behavior has never been identified in the existing laboratory literature so far (see, e.g., Fischbacher, Gächter, and Fehr (2001)). Consequently, between 4 and 17% of subjects are classified as free riders, which is significantly less than the 20 to 30% that the existing literature typically finds.²¹ We can also see from the figure that modifying the benchmark has the main effect of reclassifying some "weak reciprocators" either as "free riders" or "altruists". The proportion of "non weak" reciprocators is stable at around 41%. In all of our below analysis, we will use a benchmark of 0.2 to define the type of our subjects, since this allows for some "errors" in their contribution decisions without having a strong impact on the relative proportion of types in the population. Our results, however, are robust to varying this benchmark.²²

Figure 3 features the magnitude of developers' self-reported motives for contributing to OSS, which we elicited in our post-experimental questionnaire, together with their 95% confidence intervals. One striking result that emerges from this figure is that developers tend to disagree very much with the idea that earning money motivates their contributions, both in absolute terms (score of 0.25 out of 1, where 0 means "disagree strongly" and 1 means "agree strongly"²³) and relative to the other suggested motivations.²⁴ Developers also tend to disagree quite a bit with the idea that their contributions are motivated by the willingness to get a reputation in the OSS developers scene (score of 0.46). By contrast, the pleasure derived from learning and problem solving is the motivation that our subjects agree the most with (score of 0.76), followed by the need to solve a problem that could not be solved by proprietary software and the belief that software should not be a proprietary product (scores of 0.64 and 0.62, respectively). At the end of the day, developers tend to rank those motivations that are clearly intrinsic the highest, while rejecting those that are fully extrinsic (earning money) or have an extrinsic component (establishing one's reputation in the OSS scene).

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 OSS (see table 3). Older developers tend to contribute more to OSS across all types of contribution. Being one year older is associated with a 2% increase in the number of commits made, a 3% increase in the number of files modified a 4% increase in the number of patches submitted and a 12% increase in the number of messages posted which follow-up on some previous question asked in a project forum. On top of being vastly underrepresented in the overall population of contributors, female developers contribute significantly less than males, irrespective of the type of contribution considered. The

²¹While they are suggestive, those differences cannot be directly interpreted. Indeed, beyond the difference in subjects pool, the decision environment (i.e. the Internet) could also explain some of the observed differences in behavior (see Hergueux and Jacquemet (2014).

²²Tables available from the authors upon request.

²³The original survey questions were on a 10 points scale of agreement. We normalized the scores to be between 0 and 1 in order to ease the comparison of the magnitude of the coefficients that we report in our below regression analysis.

²⁴One striking fact is that this conclusion remains true if we restrict the sample of developers to those who are actually being paid to contribute. In this case, the "earning money" motivations reaches a 0.36 agreement score.

other socio-demographic variables are less consistently associated with developers' contribution patterns. Developers' education level is positively and significantly associated with the number of commits made and the number of follow-up forum messages posted, but not with the number of files modified and patches submitted. The coefficients on income are all insignificant and close to zero, except for the one on the number of patches submitted. This result is surprising, as it suggests that subjects' opportunity cost of time does not have a significant impact on their willingness to contribute to OSS. This conclusion is reinforced if we run the exact same regression on the sub-sample of subjects who declare not receiving any monetary compensation for their contributions, be it directly or indirectly. In this case, the coefficients on the income variable are all statistically insignificant.²⁵ Finally, moving from generally being "unwilling to take risks" to being "fully prepared to take risks" is actually associated with an 82% decrease in the number of code contributions.

With those results in mind, we now turn to our theoretical question of interest. Based on the classification of types that we derive from developers' behavior in the conditional Public Goods game – and controlling for the above survey-based variables – the next section investigates whether social motives are predictive of their willingness to contribute to OSS in the field. The last section focuses the analysis on developers' contributions of code to projects with which they are formally affiliated, and distinguishes between the determinants of the intensive and extensive margins of participation.

3.2 Motivations to contribute to OSS

Controlling for the above socio-demographic characteristics, table 4 tests for patterns of association between developers' elicited motivations and their real world contribution records, spanning the 3 distinct types of activities discussed in section 2.4. Focusing first on the motives that we elicit in the post experimental questionnaire, we find that the motivations that developers tend to rank as the most important in the survey are not the ones that are most strongly associated with their actual field contributions. Indeed, the "like to learn" and the "ideology" motives, which both achieved a high agreement ranking in the survey, are not significantly associated with developers' contribution levels in the field across all 3 types of activities. The "like to learn" motive is even significantly and *negatively* associated with developers' willingness to dedicate resources to providing solutions to other teams when they face a bug (see column (3)). This negative correlation is not surprising in the sense that contributing a patch to some other development team is costly (since the developer cannot commit his code directly to the project, but has to upload it in a separate document together with some documentation detailing how it should be implemented) and presumably yields rather low learning benefits. Similarly, the "own use" motive is not systematically associated with the number of contributions made to OSS by developers.

By contrast, the motivations that have an extrinsic component – i.e. establishing one's reputation in the OSS community and earning money – are strongly associated with higher field contri-

²⁵Table available from the authors upon request.

butions to OSS. This is a notable result, as those motivations are also the ones that are ranked the lowest in the post-experimental questionnaire. In this case, moving from a "strong disagreement" to a "strong agreement" with the idea that reputation concerns motivates contributions is associated with a 112% increase in the number of commits made and a 76% increase in the number of files modified to projects with which the developer is formally affiliated (columns (1) and (2)). It is also associated with a 151% increase in the number of patches contributed to other development teams (column (3)) and with more than a fivefold increase in the number of follow-up messages contributed to projects' forums (column (4)). Similarly, being paid to contribute to OSS, either directly or indirectly, is significantly associated with a 38% increase in the number of commits made, a 35% increase in the number of files modified and a 55% increase in the number of patches contributed to other development teams.²⁶ It is not associated with developers' level of participation in message forums, however. This result can be explained by the fact that firms lack the proper incentives to sponsor end-user support activities, since they often tend to contribute to OSS precisely in order to subsequently develop and sell customized versions of the software, together with documentation and customer support services (Lerner and Tirole (2005a)). Actually, the only motivation which is significantly associated with developers' willingness to engage in end-user support is their reputation motive. This result is consistent with Lakhani and Von Hippel (2003)'s qualitative study of the determinants of engagement in end-user support, in which they make the case that developers who are the most likely to contribute to message forums are those who want to establish themselves as topical experts within the community, and therefore seek to answer the questions that they consider as belonging to their area of expertise.

Turning our attention to the role of our experimentally elicited social motivations (and holding all other factors fixed), we see that altruism and reciprocity preferences are significant predictors of developers' contributions of code to OSS. Regardless of the type of contribution considered, free-riders systematically appear as weaker contributors than the 3 other groups. This is particularly the case for developers' willingness to engage in across team cooperation and help other developers solve the bugs that affect their own projects (column (3)). In this case weak reciprocators are estimated to contribute 180% more patches than free-riders, reciprocators 235%, and altruists 395%.²⁷ By contrast, the coefficients on the altruism and reciprocity motives, while all positive, are insignificant to explain developers' involvement in end-user support (column (4)).

Finally, the picture is more nuanced in the case of contributions of code to projects with which developers are formally affiliated. We can see from columns (1) and (2) that being a weak reciprocator or a reciprocator is significantly associated with a 43 and a 62% increase in the number of commits made, respectively, while being an altruist is significantly associated with a 65% increase

²⁶For this motive only, we do not use subjects' self-reports on whether financial rewards is a significant motive to contribute to OSS. Instead, we use the question on whether they *actually* receive direct or indirect monetary rewards for their contributions. Using the former variable yields qualitatively similar, but only marginally significant results, which also points to the fact that self reported motivations need to be treated with some care.

²⁷It is important to keep in mind that developers only contribute 1.24 patch to other development teams on average while interpreting the magnitude of those coefficients. The value-added of those external code contributions for the teams that receive them can be very high, however.

in the number of files modified. Those estimations, however, average out the extensive (i.e. in how many projects do developers seek to obtain the status of "committer" and become formally part of the development team) and intensive (i.e. how intensely do developers contribute to the projects with which they are affiliated on average) margins of code contributions. We therefore distinguish between the determinants of the intensive and extensive margins of participation in the next section.

3.3 Intensive and extensive margins of contributions

The extensive and intensive margins of developers' code contributions could have very different determinants. On the one hand, there is a relatively high entry cost to any given software project in terms of learning about its structure and functioning for an individual developer (Shah (2006)). Once that this fixed cost is paid, the marginal cost of sustaining one's contributions is lower. On the other hand, while observing individual developers' contributions of code at the project level is quite costly, obtaining committer status on a project (and, therefore, being listed as one of the members of its development team on its Sourceforge web page) is an easy to observe signal of quality for an external firm. As a result, developers who are relatively more motivated by reputation should, for instance, actively seek to gain formal affiliation with more projects.

Table 5 distinguishes between the determinants of the intensive and extensive margins of OSS participation. Column (1) explains the number of projects on which each developer has committer status (the extensive margin) while column (2) explains the number of code contributions made per project on average (the intensive margin). Columns (3) and (4) feature the same dependent variables, but focus respectively on the number of projects on which a developer is an administrator and on the number of code contributions made as a project administrator per project. Indeed, obtaining admin status on a project is an even stronger signal of commitment and quality than achieving committer status. Since the variables "nb commits" and "nb files" are heavily correlated ($corr=0.67$, $p<0.001$) and yield similar conclusions, we only report the coefficients for the "nb commits" variable in this table.

Interestingly, columns (1) and (2) of table 5 reveal that socially motivated developers all participate in significantly less projects than free-riders, but contribute significantly more intensely to each of them. Namely, being a weak reciprocator is associated with a 17% decrease in the number of projects on which the developer is a committer, but also a 73% increase in the number of commits of code made to each of them. Similarly, being a reciprocator or an altruist is associated with a 17 and a 25% decrease in the number of projects, but a 95 and a 86% increase in the number of commits made to each, respectively. This conclusion is reinforced – both in terms of magnitude and statistical significance – if we focus the analysis on the number of projects on which developers become administrators (column (3)), and on the number of contributions (3) that they make per project on which they hold admin rights (column (4)).

Turning our attention to the survey-based motivation variables, we see that developers who are relatively more motivated by getting a reputation in the open source software community

are *both* affiliated with significantly more projects (be it as a committer or as an administrator, see columns (1) and (3)) and contribute more intensely to each (see columns (2) and (4)). By contrast, being paid to contribute to OSS is significantly associated with the intensive margin of contributions, but not the extensive margin. Conversely, an interesting result relates to the "own use" motive. We can see from the table that this motive is significantly associated with the extensive margin of contributions, but not the intensive margin. This result could point to the fact that developers who contribute out of a need to fulfill some private need that cannot be met by proprietary software tend to contribute to many different projects, but to not follow-up on their initial contributions (perhaps as they stop contributing as soon as they successfully tweaked the program to meet their personal needs). Finally, consistent with the results from table 4, the "ideology" and "like to learn" motives are never significantly associated with developers' contribution patterns.

4 Developers' sorting behavior at the team level

In this section, we study whether developers tend to match assortatively into teams based on their experimentally elicited cooperative types. There is substantial experimental evidence that more cooperative types seek to match assortatively in public goods like environments, often with a significant impact on group efficiency (Page, Putterman, and Unel, 2005; Cinyabuguma, Page, and Putterman, 2005; Charness and Yang, 2008). It has typically been challenging for the literature to test for such endogenous sorting behaviors at the group level in the field, however. Our combination of experimental and field data on the community of OSS developers gives us a unique opportunity to address this question.

From our initial sample of 1,194 developers, we identify those who belong to the same development teams (i.e. those who hold commit rights on the same projects and have contributed at least one commit to those projects). We are able to identify 270 such developers, working together on 131 distinct projects. Out of the 131 teams that we identify in our sample, 93 have 2 developers, 23 have 3 developers, 12 have 4 developers and one team has 5, 6 and 8 developers, respectively.²⁸

Based upon our above classification of developers into four cooperative types, we start by describing how diverse those 131 development teams tend to be. We compute a Herfindahl index of concentration of types at the team level. We then take one minus this quantity in order to get an indicator that grows from zero to one as teams tend to be more diverse in terms of the cooperative types of their members:

$$D = 1 - \sum_{t=1}^4 p_t^2 \quad (2)$$

where p_t represents the proportion of developers who are of cooperative type t in the development team considered.

²⁸Of course, those teams can be larger in practice, as some of their members are likely not to have participated in our study.

Figure 5 features the distribution of this indicator of diversity of cooperative types across all 131 development teams. We can see that the distribution features two modes: one at zero (i.e. perfect homophily at the team level), and the other at 0.5, so that a significant fraction of teams are actually comprised of developers with different types.

In a second step, we test for homophily at the team level. For each developer i , we compute the proportion of the other members j of his team that are of his cooperative type. We then subtract from this proportion the proportion of developers who are of that particular type in the whole underlying population of developers. We therefore get, for each developer, an indicator of how much the proportion of the other members of his team who are of his own cooperative type deviates from what would have been predicted based on a random draw from the underlying population:

$$DevPropSame = \frac{\sum_{j=1}^n d_t}{n} - P_t \quad (3)$$

where d_t denotes another member j of developer i 's team who is of his cooperative type and P_t represent the proportion of developers who are of that same type in the whole underlying population of developers.

A Wilcoxon signed-rank test indicates that the *DevPropSame* indicator is statistically significantly greater than zero, therefore pointing at a tendency for developers to match assortatively by cooperative type ($z=1.74$, $p=0.082$, two-tailed). This average effect conceals some interesting type-specific heterogeneity, however. Computing the indicator separately for each type, we obtain that free-riders actually seek to be matched with developers of *different* types than their own ($z=-2.85$, $p=0.004$, two-tailed), while conditional cooperators actively seek to be matched with developers of their own type ($z=2.60$, $p=0.009$, two-tailed). This result strikingly confirms the findings of the above-mentioned experimental literature on the private provision of public goods. Indeed, this literature consistently finds that conditional cooperators tend to match assortatively and exclude free-riders from their experimental group when given the opportunity to sort endogenously. Computing the *DevPropSame* indicator for the groups of weak conditional cooperators and altruists yields statistically insignificant results ($z=-0.30$, $p=0.766$, two-tailed and $z=1.29$, $p=0.196$, two-tailed, respectively).

In table 6, we study developers' social sorting behavior in a regression framework. In column (1) we regress the *DevPropSame* variable on each developer's self-reported motivations, socio-demographic characteristics and cooperative type (taking the group of free-riders as the baseline).²⁹ Interestingly, we see that, on average, female developers seek to be matched with other developers of their own cooperative type significantly more than males, while developers who are paid to contribute seek to be matched with other developers of their own type significantly less than those who are not. The coefficients on the individual cooperative type dummies are insignificant, indicating that their taste for homophily do not statistically differ from that of the free-riders.

²⁹One limitation of this approach is that it assumes that developers of a certain type, if anything, seek to be matched with developers of their own type.

In column (2) we add a dummy variable indicating whether the developer is an administrator on the project. This is potentially important, since project administrators are the gatekeepers of their development teams and get to choose who ultimately joins to development team. The coefficient on the admin dummy is positive and (marginally) significant, meaning that project administrators do, on average, seek collaborator of their own cooperative type. Column (3) tries to make the nature of this effect more precise by estimating separate coefficients by cooperative type and administrator status (taking the free riders, non admin as the comparison group). We observe that most of the effect that we identify in column (3) actually comes from the reciprocating and altruist administrators, who are the ones who seek to be matched assortatively with other developers of their own type significantly more. Being an reciprocating administrator is estimated to cause a deviation in the proportion of the other members of the team who are reciprocators of 17 percentage points, while being an altruist administrator is estimated to cause a deviation of 30 percentage points. Since the proportion of reciprocators and altruist in the whole underlying population of developers are of 47 and 12%, respectively, the magnitude of the screening behavior of those administrators is quantitatively large.

Finally, column (4) estimates the same regression as column (1), but restricts the sample to developers who are project administrators. This allows to investigate whether certain socio-demographic characteristics or other motivational variables have a differential impact on administrators' propensity to coopt developers who are of their own type within their development teams. We see that this is actually the case, as the propensity of project administrators to seek collaborators of their own type is, on average, significantly positively associated with their ideology motive and their level of risk aversion.

At the end of the day, we find strong evidence of endogenous team-level sorting by cooperative type within our sample of OSS developers. Free-riders seek to join development teams that are comprised of different social types than their own, while strongly cooperative developers tend to match assortatively. This assortative matching effect seems to be largely driven by reciprocating and altruist project administrators, who act as gatekeepers for their development teams and seek to coopt developers of their own cooperative type.

5 Conclusion

Non monetary motives can be powerful work incentives. Especially in sectors in which it is difficult to contract upon workers' level of effort, managers of private and public organizations rely increasingly strongly on heterogenous motivations to incentivize production and innovation. In fact, the past two decades have seen the rise of a significant organizational innovation – peer production – which needs not rely on monetary incentives to elicit contributions and effort from its members, at least in certain contexts (Algan, Benkler, Morell, and Hergueux (2013)). This emerging production model is nonetheless a significant source of value creation in the most innovative and competitive sectors of information and technology.

In this paper, we rely on the open source software (OSS) peer production community to show

that social motives strongly predict developers' patterns of contributions to individual software projects. In doing so, we control for a variety of socio-demographic variables and other self-reported motivations, both intrinsic and extrinsic. We also use our combination of experimental and field data to show that developers endogenously sort by social type at the team level. While such sorting behavior is essential for organizational efficiency (Besley and Ghatak, 2005; Kosfeld and von Siemens, 2011) and public goods provision (Fischbacher and Gächter (2010)), we are the first to study this question empirically out of a laboratory context.

References

- ALGAN, Y., Y. BENKLER, M. F. MORELL, AND J. HERGUEUX (2013): "Cooperation in a Peer Production Economy: Experimental Evidence from Wikipedia," in *Workshop on Information Systems and Economics, Milan, Italy*, pp. 1–31.
- ANDERSON, S. P., J. K. GOEREE, AND C. A. HOLT (1998): "A theoretical analysis of altruism and decision error in public goods games," *Journal of Public Economics*, 70(2), 297–323.
- ANDREONI, J. (1989): "Giving with impure altruism: Applications to charity and Ricardian equivalence," *The Journal of Political Economy*, 97(6), 1447.
- (1990): "Impure altruism and donations to public goods: a theory of warm-glow giving," *The economic journal*, pp. 464–477.
- BANDIERA, O., I. BARANKAY, AND I. RASUL (2005): "Social preferences and the response to incentives: Evidence from personnel data," *The Quarterly Journal of Economics*, 120(3), 917–962.
- BARTLING, B., E. FEHR, AND H. HERZ (2013): "The intrinsic value of decision rights," Discussion paper, Working Paper Series, Department of Economics, University of Zurich.
- BARTLING, B., E. FEHR, AND K. M. SCHMIDT (2012): "Screening, Competition, and Job Design: Economic Origins of Good Jobs," *American Economic Review*, 102(2), 834–864.
- BELENZON, S., AND M. A. SCHANKERMAN (2008): "Motivation and Sorting in Open Source Software Innovation," SSRN Scholarly Paper ID 1401776, Social Science Research Network, Rochester, NY.
- BÉNABOU, R., AND J. TIROLE (2006): "Incentives and Prosocial Behavior," *The American Economic Review*, pp. 1652–1678.
- (2013): "Bonus Culture: Competitive Pay, Screening, and Multitasking," Discussion paper, National Bureau of Economic Research.
- 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.
- BERG, J., J. DICKHAUT, AND K. MCCABE (1995): "Trust, reciprocity, and social history," *Games and economic behavior*, 10(1), 122–142.
- BESLEY, T., AND M. GHATAK (2005): "Competition and incentives with motivated agents," *American economic review*, 95(3), 616–636.
- BESSEN, J. E. (2005): "Open source software: Free provision of complex public goods," Available at SSRN 588763.
- BITZER, J., W. SCHRETTL, AND P. J. SCHRÖDER (2007): "Intrinsic motivation in open source software development," *Journal of Comparative Economics*, 35(1), 160–169.
- CHARNESS, G., AND C.-L. YANG (2008): "Endogenous Group Formation and Public Goods Provision: Exclusion, Exit, Mergers, and Redemption," .

- CHAUDHURI, A. (2011): "Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature," *Experimental Economics*, 14(1), 47–83.
- CINYABUGUMA, M., T. PAGE, AND L. PUTTERMAN (2005): "Cooperation under the threat of expulsion in a public goods experiment," *Journal of Public Economics*, 89(8), 1421–1435.
- CROWSTON, K., K. WEI, J. HOWISON, AND A. WIGGINS (2012): "Free/Libre open-source software development: What we know and what we do not know," *ACM Computing Surveys (CSUR)*, 44(2), 7.
- DAVID, P. A., AND J. S. SHAPIRO (2008): "Community-based production of open-source software: What do we know about the developers who participate?," *Information Economics and Policy*, 20(4), 364–398.
- DAVID, P. A., A. WATERMAN, AND S. ARORA (2003): "FLOSS-US the free/libre/open source software survey for 2003," *Stanford Institute for Economic Policy Research, Stanford University, Stanford, CA*.
- DOHMEN, T., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP, AND G. G. WAGNER (2011): "Individual risk attitudes: Measurement, determinants, and behavioral consequences," *Journal of the European Economic Association*, 9(3), 522–550.
- DUFWENBERG, M., AND G. KIRCHSTEIGER (2004): "A theory of sequential reciprocity," *Games and economic behavior*, 47(2), 268–298.
- FALK, A., AND U. FISCHBACHER (2006): "A theory of reciprocity," *Games and Economic Behavior*, 54(2), 293–315.
- FALK, A., AND M. KOSFELD (2006): "The hidden costs of control," *The American economic review*, pp. 1611–1630.
- FEHR, E., AND C. F. CAMERER (2004): "Measuring Social Norms and Preferences Using Experimental Games: A Guide for Social Scientists," *Foundations of Human Sociality*, 1(9), 55–96.
- FISCHBACHER, U., AND S. GÄCHTER (2010): "Social preferences, beliefs, and the dynamics of free riding in public goods experiments," *The American Economic Review*, pp. 541–556.
- FISCHBACHER, U., S. GÄCHTER, AND E. FEHR (2001): "Are people conditionally cooperative? Evidence from a public goods experiment," *Economics Letters*, 71(3), 397–404.
- GAUDEUL, A. (2004): "Competition between open-source and proprietary software: the (La) TeX case study," *Working Paper, University of Toulouse*.
- GHOSH, R., R. GLOTT, B. KRIEGER, AND G. ROBLES (2002): "Free/Libre and Open Source Software: Survey and Study - FLOSS FINAL REPORT," Discussion paper.
- GHOSH, R. A. (2007): "Economic impact of open source software on innovation and the competitiveness of the Information and Communication Technologies (ICT) sector in the EU," .
- GLAZER, A. (2004): "Motivating devoted workers," *International Journal of Industrial Organization*, 22(3), 427–440.
- GREENSTEIN, S., AND F. NAGLE (2014): "Digital Dark Matter and the Economic Contribution of Apache," *Forthcoming, Research Policy*.

- HARS, A., AND S. OU (2001): "Working for free? Motivations of participating in open source projects," in *Proceedings of the 34th Annual Hawaii International Conference on System Sciences, 2001*, pp. 9 pp.—.
- HARUVY, E., F. WU, AND S. CHAKRAVARTY (2005): "Incentives for developer's contributions and product performance metrics in open source development: an empirical exploration," .
- HENKEL, J., AND M. TINS (2004): "Munich/MIT survey: development of embedded Linux," .
- HERGUEUX, J., AND N. JACQUEMET (2014): "Social preferences in the online laboratory A randomized experiment," *Forthcoming, Experimental Economics*.
- HERTEL, G., S. NIEDNER, AND S. HERRMANN (2003): "Motivation of software developers in Open Source projects: an Internet-based survey of contributors to the Linux kernel," *Research policy*, 32(7), 1159–1177.
- JULLIEN, N., AND J.-B. ZIMMERMANN (2009): "Firms' contribution to open-source software and the dominant user's skill," *European Management Review*, 6(2), 130–139.
- KOGUT, B., AND A. METIU (2001): "Open-source software development and distributed innovation," *Oxford Review of Economic Policy*, 17(2), 248–264.
- KOSFELD, M., AND F. A. VON SIEMENS (2011): "Competition, cooperation, and corporate culture," *The RAND Journal of Economics*, 42(1), 23–43.
- KREPS, D. M. (1997): "Intrinsic motivation and extrinsic incentives," *The American Economic Review*, pp. 359–364.
- KUAN, J. W. (2001): "Open source software as consumer integration into production," *Available at SSRN 259648*.
- LAKHANI, K., AND R. WOLF (2005): "Why Hackers Do What They Do: Understanding Motivation and Effort in Free/Open Source Software Projects," .
- LAKHANI, K. R., AND E. VON HIPPEL (2003): "How open source software works: free user-to-user assistance," *Research policy*, 32(6), 923–943.
- LAZER, D., A. S. PENTLAND, L. ADAMIC, S. ARAL, A. L. BARABASI, D. BREWER, N. CHRISTAKIS, N. CONTRACTOR, J. FOWLER, M. GUTMANN, ET AL. (2009): "Life in the network: the coming age of computational social science," *Science (New York, NY)*, 323(5915), 721.
- LERNER, J., AND J. TIROLE (2002): "Some simple economics of open source," *The journal of industrial economics*, 50(2), 197–234.
- (2005a): "The Economics of Technology Sharing: Open Source and Beyond," *The Journal of Economic Perspectives*, 19(2), 99–120.
- (2005b): "The Scope of Open Source Licensing," *Journal of Law, Economics, and Organization*, 21(1), 20–56.
- MAURER, S. M., AND S. SCOTCHMER (2006): "Open source software: the new intellectual property paradigm," Discussion paper, National Bureau of Economic Research.

- MOCKUS, A., R. T. FIELDING, AND J. D. HERBSLEB (2002): "Two case studies of open source software development: Apache and Mozilla," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 11(3), 309–346.
- PAGE, T., L. PUTTERMAN, AND B. UNEL (2005): "Voluntary association in public goods experiments: Reciprocity, mimicry and efficiency*," *The Economic Journal*, 115(506), 1032–1053.
- RABIN, M. (1993): "Incorporating fairness into game theory and economics," *Advances in Behavioral Economics*, p. 297.
- ROBERTS, J. A., I.-H. HANN, AND S. A. SLAUGHTER (2006): "Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the Apache projects," *Management science*, 52(7), 984–999.
- SAINT-PAUL, G. (2003): "Growth effects of nonproprietary innovation," *Journal of the European Economic Association*, 1(2-3), 429–439.
- SCHMIDT, K. M., AND M. SCHNITZER (2002): "Public Subsidies for Open Source-Some Economic Policy Issues of the Software Market," *Harv. JL & Tech.*, 16, 473.
- SHAH, S. K. (2006): "Motivation, governance, and the viability of hybrid forms in open source software development," *Management Science*, 52(7), 1000–1014.
- SHIMIZU, H., J. IIO, AND K. HIYANE (2004): "The realities of Free/Libre/Open Source Software developers in Japan and Asia," *First Monday*, 9(11).
- SILVA, J. S., AND S. TENREYRO (2006): "The log of gravity," *The Review of Economics and statistics*, 88(4), 641–658.
- SPENCE, A. M. (1974): *Market signaling: informational transfer in hiring and related screening processes*. Harvard University Press.
- WALLI, S., D. GYNN, AND B. ROTZ (2005): "The growth of open source software in organizations," *A report*.
- WOOLDRIDGE, J. M. (2010): *Econometric analysis of cross section and panel data*. MIT press.

Tables and figures

FIGURE 1: THE INSTRUCTION SCREEN OF THE PUBLIC GOODS GAME

Section 1/4 - Description	<p>In this section, groups of 4 participants (yourself and 3 other participants) are randomly formed.</p> <p>Remember: The participants who belong to your group in this section are different from those you encounter in the other sections of the study.</p> <p>At the beginning of this section, each member of the group receives \$10.</p> <p>Each member of the group must then decide how many dollars to keep for himself or herself and how many to invest in a common project.</p> <p>Each dollar invested in the common project by a member of the group yields a return of \$0.40 to each of the 4 group members (including yourself). In other words, the total amount of the contributions to the common project is multiplied by 1.6 before being evenly distributed between the 4 group members.</p> <p><u>Your earnings in dollars at the end of this section are given by:</u></p> $10 - (\text{your contribution to the common project}) + 0.4 \times (\text{total contribution to the common project})$ <p>=> The next screen gives examples...</p>
--------------------------------------	--

◀ Previous Next ▶

FIGURE 2: THE DECISION SCREEN OF THE CONDITIONAL PUBLIC GOODS GAME

Section 1/4 - Enter your decision 2/2

This is a decision screen. Once you have made your decision and clicked the "Next" button, you will not be able to go back to this screen again.

* You are now provided with a contribution table that lists **each possible average contribution that the other group members could make** (all integers between 0 and 10).

For each possible average contribution of the other group members, how much do you want to invest in the common project?

<i>If the other group members make an average contribution of:</i>	\$0	\$1	\$2	\$3	\$4	\$5	\$6	\$7	\$8	\$9	\$10
How much do you want to invest in the common project?	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

[Review description](#)

YOU CAN READ THE DESCRIPTION OF THIS SECTION AGAIN AT ANY TIME BY CLICKING HERE

← Previous

Next →

TABLE 1: CONSTRUCTING THE SAMPLE OF ELIGIBLE SUBJECTS

Nb active developers on project	License restrictiveness	Total nb of projects	Nb of projects randomly selected	Active developers randomly selected	Non active developers randomly selected
1	High	239	57	57	57
1	Moderate	57	57	57	57
1	Low	69	57	57	57
2	High	118	28	56	56
2	Moderate	28	28	56	56
2	Low	29	28	56	56
3	High	51	12	36	36
3	Moderate	12	12	36	36
3	Low	15	12	36	36
4	High	24	6	24	18
4	Moderate	13	6	24	18
4	Low	6	6	24	18
5	High	18	7	35	21
5	Moderate	9	7	35	21
5	Low	7	7	35	21
6	High	18	5	30	15
6	Moderate	5	5	30	15
6	Low	5	5	30	15
7+	.	83	83	962	249

FIGURE 3: SELF REPORTED MOTIVES FOR CONTRIBUTING TO OSS

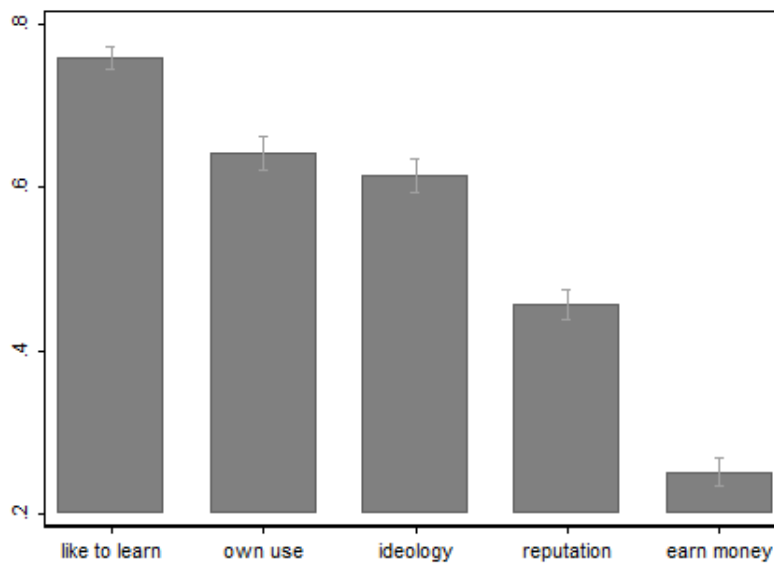


FIGURE 4: DISTRIBUTION OF COOPERATIVE TYPES IN THE CONDITIONAL PUBLIC GOOD GAME

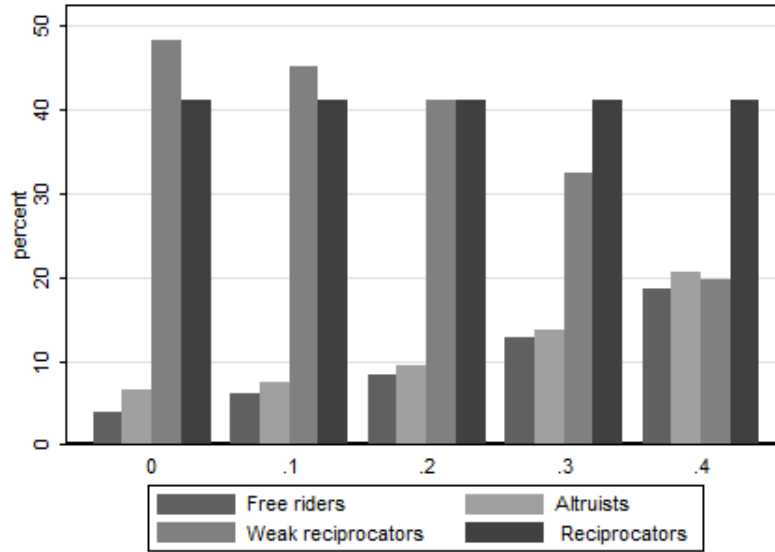


FIGURE 5: DIVERSITY OF COOPERATIVE TYPES AT THE TEAM LEVEL

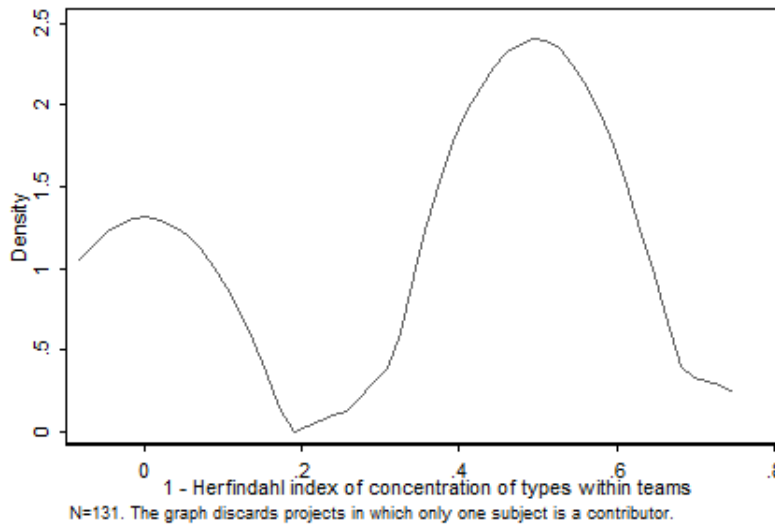


TABLE 2: DEVELOPER LEVEL DESCRIPTIVE STATISTICS

	Mean	Min	Max
Number of observations (<i>N</i>)	1194	.	.
DEPENDENT VARIABLES			
Mean – number of commits	362.77 (926.08)	0	18678
Mean – number of files added, deleted, modified and replaced	3010.17 (7197.88)	0	84459
Mean – number of patches	1.24 (5.84)	0	139
Mean – number of follow-up forum messages	25.59 (177.34)	0	2905
SOCIAL PREFERENCES MEASURES			
<i>Cooperative types – Conditional Public Goods game (N=1194, benchmark=0.2)</i>			
(i) Free riders	0.08	0	1
(ii) Weak reciprocators	0.41	0	1
(iii) Reciprocators	0.41	0	1
(iv) Altruists	0.09	0	1
<i>Altruism – Dictator (N=572)</i>			
(i) General altruism	0.38 (0.25)	0	1
(ii) Directed altruism – OSS developer	0.49 (0.24)	0	1
(iii) Directed altruism – OSS developer using GPL license	0.49 (0.25)	0	1
<i>Reciprocity – Trust (N=622)</i>			
(i) Average proportion of amount returned	0.53 (0.19)	0	1
<i>Trust in strangers – Trust (N=572)</i>			
(i) Proportion of endowment sent	0.68 (0.30)	0	1
<i>Fairness – Ultimatum (N=622)</i>			
(i) Proportion of endowment demanded	0.26 (0.23)	0	1
DEMOGRAPHIC VARIABLES			
Age	32.2 (8.42)	2	72
Proportion female	0.03 (0.16)	0	1
Degree level	5.29 (1.64)	1	8
Salary level	4.76 (2.24)	1	9
Risk aversion level	5.73 (2.40)	0	10

TABLE 3: DEMOGRAPHIC CHARACTERISTICS AND CONTRIBUTIONS TO OSS

	(1)	(2)	(3)	(4)
	Own projects nb commits	Own projects nb files	Other projects nb patches	End-user support nb messages
Age	0.02** (0.01)	0.03*** (0.01)	0.04** (0.02)	0.11*** (0.01)
Female	-0.63* (0.34)	-0.89** (0.37)	-2.02** (0.85)	-3.54*** (0.71)
Education level	0.09*** (0.04)	0.02 (0.04)	-0.05 (0.08)	0.21*** (0.07)
Income	0.02 (0.03)	-0.01 (0.03)	0.15** (0.06)	-0.07 (0.05)
Risk aversion	-0.06*** (0.02)	-0.03 (0.02)	-0.05 (0.05)	0.07* (0.04)
Observations	1103	1103	1103	1103
Pseudo R^2	0.003	0.001	0.009	0.016

Negative binomial estimates. Standard errors are reported in parentheses. Constants not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

TABLE 4: MOTIVATIONS TO CONTRIBUTE TO OSS

	(1)	(2)	(3)	(4)
	Own projects nb commits	Own projects nb files	Other projects nb patches	End-user support nb messages
Ideology	-0.10 (0.18)	-0.00 (0.20)	0.04 (0.35)	-0.46 (0.30)
Like to learn	0.25 (0.27)	-0.10 (0.28)	-1.59*** (0.52)	-0.11 (0.54)
Own use	0.24 (0.18)	0.51*** (0.20)	0.15 (0.35)	-0.21 (0.34)
Establish reputation	0.75*** (0.20)	0.56** (0.22)	0.92** (0.42)	1.88*** (0.38)
Pay	0.32*** (0.12)	0.30** (0.13)	0.44* (0.24)	0.04 (0.24)
Weak reciprocator	0.36* (0.21)	0.31 (0.23)	1.04** (0.45)	0.11 (0.41)
Reciprocator	0.48** (0.22)	0.33 (0.23)	1.21*** (0.46)	0.19 (0.41)
Altruist	0.30 (0.27)	0.50* (0.30)	1.60*** (0.57)	0.72 (0.50)
Observations	994	994	994	994
Pseudo R^2	0.006	0.003	0.020	0.023

Negative binomial estimates. Standard errors are reported in parentheses. Constants not reported. All regressions control for the socio-demographic variables reported in table 3.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

TABLE 5: INTENSIVE AND EXTENSIVE MARGIN OF CONTRIBUTIONS

	(1)	(2)	(3)	(4)
	Own projects nb projects	Own projects nb commits per project	Own projects nb projects admin	Own projects nb admin commits per project
Ideology	0.01 (0.08)	0.03 (0.15)	-0.02 (0.11)	-0.05 (0.19)
Like to learn	0.11 (0.12)	0.09 (0.22)	0.17 (0.17)	0.02 (0.27)
Own use	0.21** (0.09)	0.09 (0.15)	0.34*** (0.12)	-0.06 (0.19)
Establish reputation	0.29*** (0.09)	0.40** (0.18)	0.33*** (0.12)	0.49** (0.21)
Pay	0.09* (0.05)	0.26** (0.11)	0.05 (0.07)	0.32** (0.13)
Weak reciprocator	-0.16* (0.09)	0.55*** (0.18)	-0.34*** (0.12)	0.80*** (0.21)
Reciprocator	-0.16* (0.09)	0.67*** (0.18)	-0.23* (0.12)	0.92*** (0.21)
Altruist	-0.22* (0.12)	0.62*** (0.23)	-0.34** (0.16)	0.81*** (0.27)
Observations	994	896	994	641
Pseudo R^2	0.015	0.008	0.014	0.009

Negative binomial estimates. Standard errors are reported in parentheses. Constants not reported.

All regressions control for the socio-demographic variables reported in table 3.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

TABLE 6: MATCHING BY COOPERATIVE TYPE AT THE TEAM LEVEL

	(1)	(2)	(3)	(4)
	Deviation prop same type from random allocation	Deviation prop same type from random allocation	Deviation prop same type from random allocation	Deviation prop same type from random allocation
Age	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.01 (0.01)
Female	0.23** (0.11)	0.20* (0.11)	0.21** (0.11)	0.08 (0.18)
Education level	0.02 (0.02)	0.01 (0.02)	0.01 (0.02)	0.07 (0.04)
Income	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	-0.03 (0.03)
Risk aversion	-0.02* (0.01)	-0.02 (0.01)	-0.02 (0.01)	-0.06** (0.02)
Ideology	0.09 (0.08)	0.09 (0.08)	0.08 (0.09)	0.36** (0.18)
Like to learn	0.12 (0.13)	0.11 (0.13)	0.09 (0.14)	0.23 (0.28)
Own use	0.04 (0.09)	0.03 (0.09)	0.03 (0.09)	0.27* (0.15)
Establish reputation	0.08 (0.10)	0.09 (0.09)	0.10 (0.10)	0.22 (0.17)
Pay	-0.11* (0.06)	-0.13** (0.06)	-0.13** (0.06)	-0.20* (0.10)
Weak reciprocator	-0.00 (0.12)	0.03 (0.12)		-0.13 (0.22)
Reciprocator	-0.09 (0.13)	-0.07 (0.13)		-0.11 (0.23)
Altruist	0.10 (0.15)	0.10 (0.14)		0.13 (0.23)
Project administrator		0.11* (0.06)		
Weak reciprocator x Non project admin			0.15** (0.07)	
Reciprocator x Non project admin			0.00 (0.08)	
Altruist x Non project admin			0.20 (0.12)	
Free rider x Project admin			0.30 (0.21)	
Weak reciprocator x Project admin			0.10 (0.12)	
Reciprocator x Project admin			0.17* (0.10)	
Altruist x Project admin			0.30** (0.13)	
Observations	263	263	263	74
Adjusted R ²	0.056	0.065	0.066	0.110

OLS estimates. Robust standard errors are clustered at the project level and reported in parentheses.
 Constants not reported.

* p<0.1, ** p<0.05, *** p<0.01

Appendix

This appendix reports on the predictive power of our secondary experimental measures of reciprocity (from trustees' behavior in the Trust game) and altruism (from dictators' behavior in standard and directed Dictator games). The results are presented in table 7. The role of reciprocity as a motivational driver of OSS contributions is strongly confirmed in panel A. The coefficients on reciprocity for the aggregate number of commits made and the number of projects to which developers decide to contribute both go in the same direction, but are statistically insignificant. However, moving from no reciprocity to full reciprocity in the Trust game is positively and significantly associated with a rise in the number of commits made per project, the number of patches submitted and the number of follow-up forum messages posted.

By contrast, the experimental measures of general altruism that we derive from the standard Dictator game are not significantly associated with developers' contribution patterns (see panel B). Surprisingly, even our measures of directed altruism, which we purposefully maintained in last position in the experiment because we were worried about demand effects, are not significantly associated with higher contribution records. This is true irrespective of whether we consider altruism directed towards OSS developers (panel C) or towards OSS developers who work on projects with a GPL license (panel D). If anything, directed altruism is associated with a marginally significant *decrease* in the aggregate number of commits made (see panel C). While this result is puzzling, we interpret the fact that we get significant results using a different experimental measure of altruism (derived from the conditional Public Goods game) as invalidating the Dictator game measures of altruism rather than the underlying theory of altruism as a motive for sustaining contributions in public goods like environments.

TABLE 7: RECIPROCITY AND ALTRUISM IN THE OTHER EXPERIMENTAL GAMES

	(1)	(2)	(3)	(4)	(5)
	Own projects nb commits	Own projects nb projects	Own projects nb commits per project	Other projects nb patches	End-user support nb messages
<i>Panel A</i>					
Reciprocity - Trust	0.45 (0.43)	-0.18 (0.18)	0.81** (0.39)	2.87*** (0.97)	1.83** (0.86)
Observations	522	522	478	522	522
Pseudo R ²	0.005	0.010	0.007	0.042	0.033
<i>Panel B</i>					
General altruism - Dictator	-0.42 (0.32)	0.10 (0.17)	-0.30 (0.27)	-0.21 (0.73)	-0.38 (0.61)
Observations	472	472	418	472	472
Pseudo R ²	0.008	0.024	0.012	0.018	0.022
<i>Panel C</i>					
Directed altruism (OSS) - Dictator	-0.59* (0.35)	0.03 (0.17)	-0.46 (0.29)	0.26 (0.78)	-0.45 (0.67)
Observations	472	472	418	472	472
Pseudo R ²	0.009	0.024	0.012	0.018	0.022
<i>Panel D</i>					
Directed altruism (OSS with GPL) - Dictator	-0.30 (0.34)	0.10 (0.17)	-0.12 (0.28)	-0.29 (0.75)	0.39 (0.73)
Observations	472	472	418	472	472
Pseudo R ²	0.008	0.024	0.011	0.018	0.022

Negative binomial estimates. Standard errors are reported in parentheses. Constants not reported.

All regressions control for the socio-demographic variables reported in table 3 and for the motives elicited in the post-experimental questionnaire.

* p<0.1, ** p<0.05, *** p<0.01

TABLE 8: ROLE OF GENERAL TRUST

	(1)	(2)	(3)	(4)	(5)
	Own projects nb commits	Own projects nb projects	Own projects nb commits per project	Other projects nb patches	End-user support nb messages
Trust in strangers - Trust	-0.78*** (0.29)	-0.27* (0.14)	-0.31 (0.24)	0.58 (0.58)	-0.57 (0.57)
Observations	472	472	418	472	472
Pseudo R^2	0.009	0.026	0.012	0.018	0.023

Negative binomial estimates. Standard errors are reported in parentheses. Constants not reported.

All regressions control for the socio-demographic variables reported in table 3 and for the motives elicited in the post-experimental questionnaire.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Concluding remarks & future work

As a conclusion, I would like to emphasize three core findings that deserve further inquiry and related lines of research and that I intend to follow in some upcoming work.

First, an immediate extension of the present work is to develop a model of peer production that would formalize its sustainability conditions and establish the nature of its competitive advantage over proprietary production models. I am currently working on the construction of this theory. Some might argue that the theory of peer production should have been developed prior to running actual experiments within peer production communities. I would argue to the contrary. Indeed, while Open Source Software has received a great deal of attention from economists, most of the existing literature on this topic is already theoretical. The models in this literature, on top of being specifically geared towards OSS as opposed to the broader concept of peer production, systematically make implicit assumptions about the nature of contributors' motivations that could not be supported by empirical evidence so far. Those implicit assumptions nonetheless largely determine the circumstances under which OSS is shown to dominate proprietary software and to enhance social welfare.¹

Furthermore, this dissertation demonstrates that many aspects of peer production can be understood within the frame of existing economic theory (e.g. public goods provision, labor economics) and that peer production communities can be powerful field sites to apply and test this theory. As a result, only in a second step does it make sense to develop a dedicated theory of peer production that would (i) be based on sound motivational assumptions and (ii) relate to existing economic theory in a tractable and parsimonious way.

Second, while the bulk of this dissertation focuses on identifying patterns of association between the magnitude (or, alternatively, the presence) of certain experimentally-elicited motives and field outcomes at the individual level, the results also suggest that those individuals who self select into contributing to peer production models may exhibit unusually high *levels* of prosocial preferences. This is easily seen by noticing that approximately 13% of Wikipedia contributors and 10% of Open Source Software developers in the sample can be classified as unconditional cooperators – i.e. “altruists” – in the conditional Public Goods game. This is a surprising pattern of behavior, which had never been documented in the existing laboratory literature.

Of course, it would be tricky to assess the importance of this self-selection effect by directly comparing the magnitude of the elicited social preferences within this population of subjects to that of a more traditional population of undergraduate students. One reason is that the Internet-based elicitation method could account for part of the likely point-estimate differences. Another reason is that it is not obvious that the relevant comparison group here should be a population of undergraduate students. Rather, it would be more informative to recruit a representative sample of the US population as the relevant comparison group, and elicit their social preferences using the exact same Internet-based methodology. Such an inquiry would, moreover, be of scientific interest in its own right.

¹ See Chapter 3, footnotes 3 and 4 for more details on this point.

One related point is that developers who decide to contribute to Open Source Software development platforms are likely to forgo significant monetary payoffs, which they could derive from participating in similar – but non collaborative – online software development platforms. TopCoder, for instance, is a centralized online platform in which developers can compete to create components and applications that third parties (e.g. firms, governments) have contracted for. The developer who submits the best proposal wins a money prize determined in advance. It would therefore be interesting to run the same online experiment on the population of software developers registered with TopCoder in order to (i) study how cooperativeness relates to performance in this highly individualistic and competitive environment and (ii) compare the magnitude of social preferences in the population of cooperative OSS developers and in the population of competitive (and presumably more extrinsically motivated) TopCoder developers.

My last point will be a methodological one. The fact that most of the experimental measures used in this dissertation have significant predictive power over economic agents' decisions in the field is certainly a great validation of the "ecological validity" of those experimental protocols. It also illustrates the usefulness of coupling experimental methods with computational social science techniques in order to help bridge the gap between experimental and observational data in economic research.

One striking result from this dissertation, however, is that behavior in the Dictator game – which is certainly the workhorse for studying altruistic preferences in the literature – appears to (i) be significantly more sensitive to context and framing than behavior in any other game (see chapter 1) and (ii) never significantly correlates with subjects' field behavior (see chapters 2 and 3). This second observation could be interpreted as invalidating the theory based on altruistic preferences. It could also be interpreted as challenging the ecological validity of this decision problem, which, if it corresponds well to the theory that it was originally supposed to test, may not be well suited to learning about subjects' behavior outside of the lab. The fact that an alternative measure of altruism constructed from subjects' behavior in the conditional Public Goods game does significantly correlate with their field behavior suggests that the second interpretation might be correct. This point certainly deserves further scrutiny.

**TRADUCTION DE L'INTRODUCTION ET DE LA
CONCLUSION DE LA THESE EN FRANCAIS**

Introduction générale & projet de recherche

Nous vivons une part croissante de nos vies sur Internet. Des forums de discussion et les sites de réseaux sociaux (Facebook par exemple) à ligne massivement multijoueur jeu de rôle (par exemple World of Warcraft), des sites d'enchères en ligne (par exemple eBay) aux marchés du travail en ligne (par exemple odesk), le cyberspace est maintenant une partie importante de le monde social et économique "réelle".

Dans le domaine de l'économie expérimentale, il a été un long temps puisque les chercheurs ont fait appel à l'élaboration de la «laboratoire en ligne». Il ya trois principales raisons de l'intérêt de l'expérimentation expérimentateurs en ligne.

Tout d'abord, l'Internet est intéressante parce qu'elle permet d'atteindre une plus grande diversité d'échantillons, de recruter de plus en plus de piscines parfois soumis représentant et même effectuer des expériences sociales interculturelles en temps réel à un coût abordable. Dans un article fondateur, Henrich et al. (2010) mettent en garde contre les spécialistes du comportement «courant trop compter sur des données recueillies auprès de populations majoritairement des étudiants de premier cycle et de recommander un effort important dans l'élargissement de la base de l'échantillon:

«Même si nous ne sommes certainement pas le premier à se soucier de la représentativité des échantillons de premier cycle qui prévalent dans les sciences du comportement, nos efforts pour compiler un cas empirique ont révélé une situation encore plus alarmante que précédemment reconnu. L'échantillon de premier cycle occidentales contemporaines qui accable si notre base de données est non seulement un échantillon restreint extraordinaire de l'humanité; il est souvent une valeur aberrante distincte vis-à-vis des autres échantillons globaux »(p. 82).

Tableau 1 documente le nombre d'utilisateurs d'Internet, le taux de pénétration d'Internet et le taux de croissance de la population des utilisateurs d'Internet par les régions du monde à partir de Juin 2012. L'image semble présenter des arguments convaincants pour l'utilisation de l'Internet comme un moyen de mener des expériences avec des échantillons importants et diversifiés: il est maintenant possible pour atteindre 78,6% de la population nord-américaine à travers l'Internet avec une relative facilité, et alors que seulement 15,6% de la population africaine ne peut actuellement être atteint grâce à cette méthode, la croissance exponentielle de son base d'utilisateurs, il sera bientôt un outil intéressant pour la réalisation d'expériences dans le monde en développement.

Deuxièmement, comme nous passons une partie importante de notre temps dans le réseau, l'Internet devient un champ expérimental de premier plan de la recherche dans son propre droit (voir Bainbridge 2007). Il est donc beaucoup de sens à mener des expériences directement sur Internet pour comprendre les différents types d'activités économiques et sociales que les gens se livrent à ligne. Un nombre croissant de documents de l'économie se servent d'expériences de terrain (Levitt et List 2009). Ceux-ci sont attrayants parce que, contrairement à des expériences de laboratoire, ils "creat [e] un contexte qui est similaire à celle dans laquelle les agents économiques opèrent" (Loewenstein, 1999). En conséquence, des expériences de terrain sont pensés pour être plus à l'extérieur valable, en ce sens que leurs résultats se généralisent plus facilement aux phénomènes du "monde réel" à partir de laquelle les chercheurs cherchent à apprendre. Techniquement parlant, l'attrait de l'Internet pour mener des expériences économiques sur le terrain est encore renforcée par le fait qu'il apparaît comme un environnement naturel pour les sujets d'interagir avec des degrés d'anonymat variable et, si nécessaire, sans répétition.

Enfin, l'expérience conduite sur Internet permet au chercheur de raconter ses résultats expérimentaux à dossiers détaillés du monde réel du comportement sur le terrain de ses sujets, donc d'aider à combler le fossé entre les données expérimentales et d'observation. En effet, dans de nombreux cas, il est possible d'extraire de l'Internet une mine d'informations valides de l'extérieur sur les décisions individuelles et de suivre leur comportement. Selon les hypothèses retenues ou sur la question de recherche posée par le chercheur, cette caractéristique de l'Internet peut permettre soit de soulager la tension traditionnelle entre validité interne et externe dans les expériences économiques, ou pour évaluer la validité écologique de procédures expérimentales qui sont utilisés largement dans le laboratoire pour tester la théorie économique. Cette caractéristique de l'Internet, il est un domaine de recherche très prometteur pour les économistes et autres spécialistes des sciences sociales. Dans un article fondateur, Lazer et al. (2009) appellent pour des investissements massifs dans la collecte et l'analyse de données largement inexploitées dans le monde numérique pour les sciences sociales systématique:

"La capacité de recueillir et d'analyser des quantités massives de données a indubitablement transformé des domaines tels que la biologie et la physique. L'émergence d'une telle «science sociale de calcul" guidée par les données a été beaucoup plus lente, en grande partie menée par quelques informaticiens, physiciens et spécialistes des sciences sociales intrépides. Si l'on regarde les grandes revues disciplinaires en économie, sociologie et les sciences politiques, il y aurait peu de preuves d'une science sociale de calcul émergents engagés dans la modélisation quantitative de ces nouveaux types de traces numériques. »(P.2)

Malgré ces caractéristiques intéressantes, cependant, et même si l'Internet est désormais une technologie bien établie dans le monde développé, le développement du "laboratoire en

ligne" reste encore à ses balbutiements. Pour être sûr, il ya des défis méthodologiques et pratiques importants associés à l'expérimentation en ligne. La première contribution de cette thèse sera donc d'ordre méthodologique. Elle fera rapport sur la construction et l'évaluation d'un nouvel outil pour mener des expériences en ligne.

Un des aspects les plus intéressants de l'Internet est que la diffusion et l'interconnexion des ordinateurs personnels a considérablement réduit les coûts de communication et d'information accroissement des flux entre les agents économiques. Cette évolution technologique a conduit à une augmentation de l'efficacité du marché et, sans doute, le bien-être social. ODesk - un marché en ligne pour main-d'œuvre qualifiée dans laquelle les programmeurs peuvent soumissionner pour exécuter des contrats modulaires postées par les clients - est un bon exemple de l'augmentation de l'efficacité du marché, une réduction significative des coûts de communication peuvent favoriser.

Au-delà des gains d'efficience du marché, cependant, l'Internet a également favorisé l'émergence d'un modèle qualitativement distinct de la production - la production par les pairs - aux côtés des entreprises, des marchés et de l'Etat, qui est principalement basé sur des contributions volontaires et la collaboration à grande échelle (Benkler 2002, 2006). Production par les pairs représente une innovation organisationnelle importante dans laquelle les agents volontairement auto-attribuent les tâches et coordonner avec succès vers la fourniture de biens publics mondiaux, en l'absence de signaux de prix et sans aucune règle de conception pré-spécifié ou de leadership formel. Au cours des deux dernières décennies, ce modèle d'organisation émergente a été une source importante de création de valeur dans les secteurs les plus innovants de l'information et de la technologie. Production par les pairs est donc intéressant d'étudier en soi. Wikipedia, le premier champ expérimental de cette thèse, accueille actuellement plus de 25 millions d'articles dans 285 librement utilisables langues. Sa valeur informative révélé semble être énorme pour la société. Il est le 5e site le plus visité sur Internet, recevoir plus de 500 millions de visiteurs uniques par mois dans le monde entier. 60% des médecins européens déclarent en utilisant Wikipedia à des fins professionnelles, et une première évaluation de la qualité de ses entrées scientifiques effectivement trouvé qu'ils étaient pratiquement impossibles à distinguer de ceux de l'encyclopédie Britannica éditée professionnellement (Giles 2005).

Plates-formes de développement de logiciels open source, le second champ expérimental de cette thèse, comportent environ 800.000 développeurs du monde entier. Logiciel Open Source (OSS) est une source majeure d'innovation, qui est responsable de la plupart des services publics de base sur lequel les pistes d'Internet (par exemple, le serveur Web Apache, Sendmail, le BIND logiciel de gestion Domain Name System), langages de programmation (par exemple, Python, Perl) et les environnements de programmation (par exemple Eclipse). OSS également en concurrence avec succès avec plusieurs de ses homologues fondées sur l'entreprise dans le domaine des systèmes d'entreprise (par exemple Linux, qui a été adopté

par un certain nombre de grandes entreprises commerciales telles que IBM, Apple et Sun pour sa fiabilité relative, la résistance aux attaques de virus, et bug vitesse de correction) et les applications de l'utilisateur final (par exemple Android, OpenOffice, VLC Media Player, Mozilla Firefox). La valeur économique générée par l'OSS est estimée à importante. En effet, Walli, Gynn, et Rotz (2005) rapportent que 87% des entreprises américaines comptent maintenant sur les logiciels libres pour certaines de leurs activités quotidiennes. Ghosh (2007) estime que le coût de recréer le code source ouvert existant à 12 milliards d'euros, tandis que Greenstein et Nagle (2014) estiment que la valeur du serveur Web Apache seul à se situer entre 2 et 12 milliards de dollars.

Mais même au-delà de sa relative nouveauté et l'importance économique, la production par les pairs représente également une occasion de jeter un nouvel éclairage sur un certain nombre de questions importantes, mais difficiles notamment dans la littérature. Comme Benkler (2013) notes:

"Les implications de la production par les pairs sont plus larges que l'impact économique direct de la pratique. Au-delà de l'ampleur de ses effets sur l'innovation et la production de connaissances dans l'économie en réseau et la participation dans la société en réseau, le succès de la production par les pairs et la coopération en ligne a plusieurs conséquences pour l'économie en général. Elle exige que nous affinons nos idées sur la motivation ou des incitations; il recalibrer les rôles de propriété et de contrat [...] dans les domaines de la production et de l'innovation de la connaissance dépend indispensables à la croissance; et il nécessite des adaptations à la théorie de l'entreprise ". (p. 2)

Compte tenu de la nature non conventionnelle de la plupart des incitations au travail en jeu dans des environnements de production par les pairs, cette thèse exploite le contexte de la production par les pairs à la recherche de l'impact des préférences économiques non standards sur la fourniture de biens publics, d'étudier leur rôle en tant que les incitations au travail, et d'évaluer leurs conséquences en termes d'économie d'organisation.

Bien que la production par les pairs comme un phénomène ne sont pas totalement nouveau pour les économistes, cette thèse est le premier à introduire le concept et prendre Wikipedia et logiciels Open Source comme instanciations particulières d'un modèle d'organisation émergente, qui devrait être durable lorsque (i) le processus de production est hautement modulaire (ie il peut être divisé en sous-tâches petites et indépendantes), (ii) les modules sont intrinsèquement motivant et (iii) la production ne nécessite pas d'intrants élevés en capital (comme les communautés de bénévoles semblent bien organiser collectivement le travail, mais pas gestion du capital).

Parce que la production par les pairs, lorsqu'elles sont durables, élimine essentiellement les coûts de passation de marchés et les problèmes d'agence (comme contributeurs sont

intrinsèquement motivés pour travailler), et accélère la découverte par la divulgation automatique, il peut bénéficier d'un avantage concurrentiel par rapport aux modèles de production exclusifs où le processus de production nécessite la participation d'un très diversifiée et main-d'œuvre qualifiée.

Alternativement, il pourrait également être appliquée avec succès à un processus d'innovation cumulative qui a besoin de faire appel à un large éventail de compétences (voir, par exemple, Maurer & Scotchmer (2006) pour l'exemple de la découverte de médicaments et l'innovation dans le secteur de la biotechnologie). Une des conséquences de l'argument ci-dessus est que, dans un cadre d'innovation cumulative, la production par les pairs bénéficieront d'un avantage comparatif important pour le développement d'innovations de rupture, car il n'y a pas de structure de gestion hiérarchique de décourager les directions qui sont très incertaines et / ou non immédiatement rentables.

Feuille de route de la thèse

La première contribution de cette thèse est d'ordre méthodologique. Chapitre 1 rapports sur le développement d'une plate-forme d'expérimentation en ligne spécialement conçu pour renforcer la validité interne des décisions induites sur Internet. La plate-forme fournit des contrôles sur la plupart des facteurs de confusion qui pourraient empêcher les expérimentateurs de mener des expériences sur Internet. En particulier, il (i) contrôle des différences de temps de réponse, (ii) traite des questions d'attrition sélective, la concentration et la distraction et (iii) fournit autant de contrôle que possible sur les croyances des sujets en ce qui concerne les instructions expérimentales. La méthodologie est appliquée à l'explicitation des préférences sociales et des risques au sein d'un échantillon de sujets de laboratoire traditionnels. Etant donné que l'interface de décision est utilisable tel quel dans le laboratoire (par le biais d'un navigateur Internet), les sujets sont répartis au hasard soit un "laboratoire" ou un état "Internet" (à savoir à la maison) à des fins de comparaison. Dans l'ensemble, en utilisant le même sujet piscine, la même interface de décision et les mêmes enjeux monétaires, la comparaison conclut en faveur de la fiabilité des comportements suscité par Internet, selon les commandes supplémentaires de la conception.

Le chapitre 2 de cette thèse prend Wikipedia comme un exemple paradigmatique de la production par les pairs, dont les motivations extrinsèques ne jouent aucun rôle dans le façonnement du comportement des contributeurs. Elle repose sur ce terrain d'étude idéal pour fournir le premier test complet sur le terrain de la part relative de chaque classe de motif sociale que la théorie économique a mis en avant pour expliquer la volonté des gens à soutenir la coopération dans les biens publics comme les environnements. En effet, les modèles basés sur les préférences de l'altruisme, la réciprocité et de l'image sociale

concurrent ont été testé en laboratoire, mais il est très peu de preuves sur le terrain à ceux qui de question plus dans des contextes économiquement pertinentes.

Ce faisant, ce chapitre illustre les avantages de la recherche qui découlent du couplage de méthodes expérimentales et informatiques dans un contexte en ligne. Bien qu'il soit possible d'obtenir de manière fiable l'altruisme et de réciprocité des préférences des sujets avec une expérience en ligne, il a généralement été un défi pour obtenir des motifs de l'image sociale expérimentale (d'autant plus dans un mode hors contexte). Pour atteindre cet objectif, le chapitre se fonde sur la richesse des données d'observation disponibles à partir de Wikipedia sur le comportement des contributeurs pour construire des mesures individuelles de préférences révélées pour l'image sociale au sein de la communauté Wikipedia. Ces mesures reposent sur la propension »(i) des contributeurs de publier plus ou moins d'informations sur eux-mêmes sur leur page Wikipedia de l'utilisateur (qui est d'aucune utilité directe pour contribuer efficacement à l'encyclopédie) et (ii) leur propension à afficher bien en évidence les signes de la vie sociale reconnaissance qu'ils ont reçu d'autres contributeurs à l'ensemble de la communauté. Il est alors possible de relier les préférences des sujets pour dossiers détaillés du monde réel de leurs contributions au projet Wikipedia, qui sont séparément extractible à partir du site.

Basé sur un échantillon représentatif de 850 contributeurs de Wikipedia, le chapitre rend compte que la réciprocité et de l'image sociale sont les deux motivations fortes pour soutenir la coopération dans ce biens publics tels que l'environnement, tandis que l'altruisme semble jouer un rôle moins important. Ce résultat confirme étonnamment les conclusions de la littérature de laboratoire existant sur la fourniture privée de biens publics. Un résultat important est que dans l'ensemble des spécifications, de la réciprocité et de l'image sociale apparaît systématiquement les conducteurs de motivation comme substituables plutôt que des complémentaires. Le chapitre décrit également les modèles spécifiques de contributions des administrateurs de Wikipédia, un groupe spécifique de contributeurs qui auto-sélectionné en jouant un rôle de maintien de l'ordre au sein de la communauté et sont notamment en charge de traiter avec les utilisateurs perturbateurs. Bien que la confiance dans les étrangers anonymes (telle que mesurée par un jeu de confiance standard) est sans rapport avec les niveaux de contribution entre les contributeurs réguliers, les estimations montrent que les administrateurs qui sont moins confiants sont beaucoup plus actif et plus susceptibles d'exercer leurs droits de maintien de l'ordre.

Le chapitre 3 de la thèse apporte une contribution à l'économie du travail en apportant économie expérimentale au sein des organisations productives du monde réel. Le chapitre se concentre sur la communauté des développeurs de logiciels libres (OSS) et provoque leurs motivations sociales avec une expérience en ligne. Comme dans le cas de Wikipedia, les activités en ligne et les contributions des développeurs individuels sont traitables à un

niveau très détaillé, ce qui est rarement le cas dans les environnements traditionnels de l'entreprise.

La combinaison des données expérimentales et de terrain sur la communauté des développeurs de logiciels libres offre l'occasion d'aborder un certain nombre de questions de longue date mais réputés difficiles dans la littérature. Il permet notamment d'étudier comment motivations hétérogène affecte l'étendue et la nature des contributions individuelles aux efforts du groupe de haut niveau. Surtout, la combinaison de données expérimentales et de terrain permet de déterminer le comportement de tri par type endogène sociale au niveau de l'équipe. En effet, alors qu'il existe des preuves théoriques et expérimentales important que les types plus coopératives cherchent à correspondre homogames au sein des groupes et des organisations afin de maintenir des niveaux élevés de coopération (qui a généralement un impact significatif sur l'efficacité), la difficulté à recueillir les données nécessaires a empêché la à partir de la littérature existante pour tester de tels comportements de tri endogènes dans le domaine. Enfin, l'approche permet également d'évaluer l'impact de la composition de l'équipe sur la probabilité de succès au niveau du projet.

Deux caractéristiques de la communauté du logiciel libre sont indispensables pour effectuer les tests ci-dessus et la distinguer de la communauté de contributeurs de Wikipédia: (i) de nombreuses équipes de développement auto-formé travaillant sur des projets de logiciels distincts coexistent simultanément et (ii) environ la moitié des développeurs qui contribuent OSS à tirer effectivement un paiement monétaire de leurs contributions.

Basé sur un échantillon stratifié de 1 194 développeurs de logiciels libres, le chapitre indique que les motivations sociales prédire les contributions des promoteurs aussi fortement que extrinsèques, mais ont un impact différent sur la nature de la participation. Socialement développeurs motivés, ont tendance à rejoindre les équipes de moins de développement (ils ont une marge plus faible étendue de la participation) mais contribuer significativement plus à chaque (avec une association positive avec l'ensemble des contributions). Le chapitre présente également des preuves solides d'un développement endogène tri au niveau de l'équipe par type de collaboration au sein de l'échantillon de développeurs de logiciels libres. Free-riders cherchent à rejoindre les équipes de développement qui sont composées de différents types sociaux que leur propre, tandis que les développeurs ont tendance à fortement coopération correspondre homogames. Ce résultat constitue la première validation de champ d'un résultat de laboratoire conforme à la littérature expérimentale sur la fourniture privée de biens publics. L'effet correspondant assortatif semble être en grande partie tirée par la réciprocité et les administrateurs du projet altruiste, qui obtiennent en fin de compte de choisir qui rejoint à l'équipe de développement et donc agissent comme les gardiens de leurs équipes en cherchant à coopter les développeurs de leur propre type coopératif.

La figure 1 résume la logique d'ensemble de ce travail, ainsi que la contribution de cette thèse à la littérature. Chapitre 1 développe et méthodologique évalue la fiabilité d'un nouvel outil pour mener des expériences économiques en ligne. Chapitre 2 et 3 document, la montée d'un nouveau modèle important pour l'organisation de la production - production pairs - dans lequel le travail des agents volontairement auto-affectation et coordonner avec succès vers la fourniture de biens publics mondiaux sans compter nécessairement sur les incitations monétaires. Ces chapitres utilisent ensuite le contexte de la production par les pairs et reposent sur une combinaison de méthodes expérimentales et informatiques pour respectivement (i) fournir le premier test complet sur le terrain de la théorie de la fourniture privée de biens publics, (ii) d'étudier l'importance des préférences sociales comme motifs de travail au sein des organisations productives du monde réel et (iii) un rapport de la première preuve de domaine du comportement de tri endogène des agents économiques au sein des équipes de production en fonction de leurs types de coopération.

Remarques de conclusion & travail futur

En conclusion, je voudrais souligner trois conclusions fondamentales qui méritent une nouvelle enquête et lignes connexes de recherche et que je compte suivre dans certains travaux à venir.

Tout d'abord, une extension immédiate de ce travail est de développer un modèle de production par les pairs qui formaliser ses conditions de durabilité et d'établir la nature de son avantage concurrentiel par rapport aux modèles de production exclusifs. Je travaille actuellement sur la construction de cette théorie. Certains pourraient soutenir que la théorie de la production par les pairs aurait dû être élaborée avant d'exécuter des expériences réelles au sein de communautés de production par les pairs. Je dirais le contraire. En effet, alors que logiciel Open Source a reçu beaucoup d'attention de la part des économistes, la plupart de la littérature existante sur ce sujet est déjà théorique. Les modèles de cette littérature, en plus d'être spécifiquement orientée vers l'OSS, par opposition à la notion plus large de la production par les pairs, font systématiquement des hypothèses implicites sur la nature des motivations des contributeurs qui ne pouvait être étayée par des preuves empiriques à ce jour. Ces hypothèses implicites néanmoins largement déterminer les circonstances dans lesquelles l'OSS est indiqué à dominer les logiciels propriétaires et d'améliorer le bien-être social.

En outre, cette thèse montre que de nombreux aspects de la production par les pairs peuvent être comprises dans le cadre de la théorie économique existante (par exemple, la fourniture de biens publics, l'économie du travail) et que les communautés de production par les pairs peuvent être des sites de terrain puissants d'appliquer et de tester cette théorie. En conséquence, seulement dans un second temps cela fait-il sens à développer une théorie dédiée de production par les pairs qui (i) sont fondés sur des hypothèses de motivation sonores et (ii) se rapportent à la théorie économique existante d'une manière souple et parcimonieux.

Deuxièmement, alors que la majeure partie de cette thèse porte sur l'identification des modes d'association entre l'ampleur (ou, à défaut, la présence) de certains motifs expérimentalement induites et les résultats sur le terrain au niveau individuel, les résultats suggèrent également que les personnes qui se sélectionner en contribuer à scruter les modèles de production peut présenter des niveaux anormalement élevés de préférences prosociales. Ceci est facilement visible en remarquant que environ 13% des contributeurs de Wikipédia et 10% des développeurs de logiciels libres dans l'échantillon peuvent être classés comme coopérateurs inconditionnels --à-dire «altruistes» - dans le jeu de biens publics

conditionnelle. Ceci est un modèle de comportement surprenant, qui n'a jamais été documenté dans la littérature existante en laboratoire.

Bien sûr, il serait difficile d'évaluer l'importance de cet effet d'auto-sélection en comparant directement l'ampleur des préférences sociales suscitées au sein de cette population de sujets à celui d'une population plus traditionnelle des étudiants de premier cycle. Une des raisons est que la méthode de sollicitation sur Internet pourrait expliquer en partie les différences estimation ponctuelle probables. Une autre raison est qu'il est pas évident que le groupe de comparaison pertinent ici doit être une population d'étudiants de premier cycle. Au contraire, il serait plus instructif de recruter un échantillon représentatif de la population des États-Unis que le groupe de comparaison pertinent, et de connaître leurs préférences sociales en utilisant la même méthodologie basée sur Internet exacte. Une telle enquête serait, en outre, être d'un intérêt scientifique à part entière.

Un autre point est que les développeurs qui décident de contribuer à ouvrir les plates-formes de développement de logiciels libres sont susceptibles de renoncer gains monétaires importants, qu'ils pourraient tirer de la participation à la même - mais non de collaboration - les plates-formes en ligne de développement de logiciels. TopCoder, par exemple, est une plate-forme en ligne centralisé dans lequel les développeurs peuvent concourir à créer des composants et applications tiers (par exemple les entreprises, les gouvernements) ont contractés. Le développeur qui soumet la meilleure proposition remporte un prix d'argent déterminée à l'avance. Il serait donc intéressant d'exécuter la même expérience en ligne sur la population de développeurs de logiciels enregistrés avec TopCoder afin de (i) l'étude comment esprit de coopération concerne la performance dans cet environnement hautement individualiste et compétitive et (ii) de comparer l'ampleur des préférences sociales dans la population de développeurs de logiciels libres et de coopération dans la population de développeurs TopCoder compétitifs (et probablement plus motivés extrinsèquement).

Mon dernier point sera d'ordre méthodologique. Le fait que la plupart des mesures expérimentales utilisées dans cette thèse ont un pouvoir prédictif significatif sur les décisions des agents économiques dans le domaine est certainement un grand validation de la «validité écologique» de ces protocoles expérimentaux. Il illustre également l'utilité de couplage des méthodes expérimentales avec des techniques de sciences sociales de calcul afin d'aider à combler l'écart entre les données expérimentales et d'observation en matière de recherche économique.

Un résultat frappant de cette thèse, cependant, est que le comportement dans le jeu du dictateur - qui est certainement le cheval de bataille pour l'étude des préférences altruistes dans la littérature - semble (i) être beaucoup plus sensibles au contexte et le cadrage de

comportement dans aucun autre jeu (voir chapitre 1) et (ii) en corrélation jamais de manière significative avec le comportement du champ de sujets (voir les chapitres 2 et 3). Cette seconde observation pourrait être interprétée comme invalidant la théorie basée sur les préférences altruistes. Il pourrait également être interprété comme une contestation de la validité écologique de ce problème de décision, qui, si elle correspond bien à la théorie selon laquelle il a été à l'origine censé tester, peut ne pas être bien adapté à l'apprentissage sur le comportement de sujets à l'extérieur du laboratoire. Le fait qu'une mesure alternative de l'altruisme construit à partir du comportement des sujets dans le jeu des biens publics conditionnelle ne corrélation significative avec leur comportement sur le terrain montrent que la seconde interprétation est peut-être correcte. Ce point mérite certainement un examen plus approfondi.

Références

- Aghion, P., Dewatripont, M. & Stein, J., 2005. *Academia, the private sector, and the process of innovation*, National Bureau of Economic Research.
- 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., 2006. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*, Yale University Press.
- 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.
- Ghosh, R. A., 2007. *Economic impact of open source software on innovation and the competitiveness of the Information and Communication Technologies (ICT) sector in the EU*. Maastricht: UNU-MERIT.
- Giles, J., 2005. Internet encyclopaedias go head to head. *Nature*, 438(7070), pp.900–901.
- Greenstein, S., & Nagle, F., 2014. Digital Dark Matter and the Economic Contribution of Apache. *Forthcoming, Research Policy*.
- Henrich, J., Heine, S.J. & Norenzayan, A., 2010. The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), pp.61–83.
- Lazer, D. et al., 2009. Life in the network: the coming age of computational social science. *Science*, 323(5915), pp.721–723.
- Lerner, J. & Tirole, J., 2002. Some Simple Economics of Open Source. *The Journal of Industrial Economics*, 50(2), pp.197–234.
- Lerner, J. & Tirole, J., 2005. The Economics of Technology Sharing: Open Source and Beyond. *The Journal of Economic Perspectives*, 19(2), pp.99–120.
- Levitt, S.D. & List, J.A., 2009. Field experiments in economics: The past, the present, and the future. *European Economic Review*, 53(1), pp.1–18.
- Loewenstein, G., 1999. Experimental Economics From the Vantage-point of Behavioural Economics. *The Economic Journal*, 109(453), pp.25–34.
- Maurer, S.M. & Scotchmer, S., 2006. *Open Source Software: The New Intellectual Property Paradigm*, National Bureau of Economic Research.
- Walli, S., Gynn, D., & Rotz, B., 2005. *The growth of open source software in organizations*. Report.
- Williams, H.L., 2013. Intellectual Property Rights and Innovation: Evidence from the Human Genome. *Journal of Political Economy*, 121(1), pp.1–27.

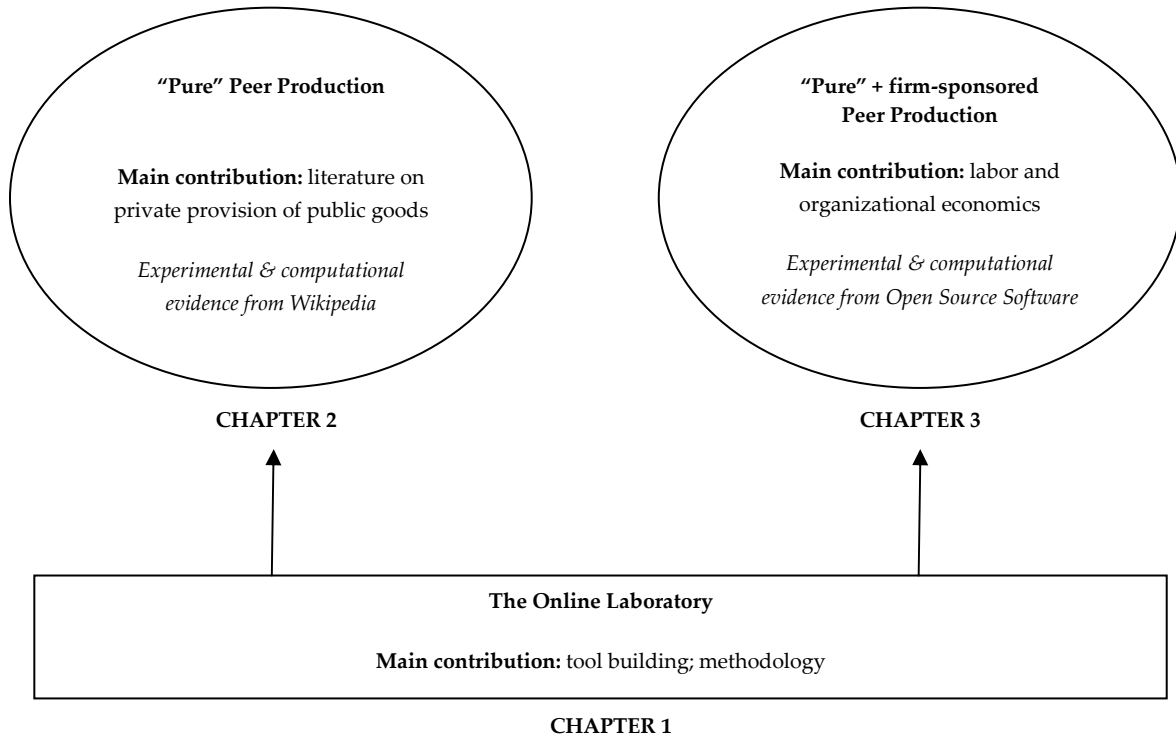
Tableaux et figures

Table 1. Utilisateurs d'Internet dans le monde et statistiques de croissance et de pénétration

World region	Total Population (June 30, 2012)	nb of Internet users (Dec. 31, 2000)	nb of Internet users (June 30, 2012)	Internet penetration rate in % (June 30, 2012)	nb of Internet users growth rate in % (2000-2012)
Africa	1,073,380,925	4,514,400	167,335,676	15.6	3,606.7
Asia	3,922,066,987	114,304,000	1,076,681,059	27.5	841.9
Europe	820,918,446	105,096,093	518,512,109	63.2	393.4
Middle East	223,608,203	3,284,800	90,000,455	40.2	2,639.9
North America	348,280,154	108,096,800	273,785,413	78.6	153.3
Latin America	593,688,638	18,068,919	254,915,745	42.9	1,310.8
Oceania	35,903,569	7,620,480	24,287,919	67.6	218.7
World total	7,017,846,922	360,985,492	2,405,518,376	34.3	566.4

Source: www.internetworldstats.com

Figure 1. La logique de la thèse



Online Cooperation and Peer Production

Abstract

Internet is a very attractive technology for the implementation of experiments. It allows to obtain larger and more diverse samples and gives the researcher the opportunity to extract from the Internet a wealth of field data that document the real-world decisions and behavior of his subjects. Notwithstanding those appealing features, the development of the “online laboratory” remains in its infancy, mainly because of the threats to validity and important practical challenges typically associated with online experimentation.

More than a potentially powerful medium to run experiments, the Internet is also a very promising field of economic research. Over the past 20 years, its diffusion has significantly reduced communication costs and increased information flows between economic agents. This technological change has notably fostered the emergence of a new production model – peer production – which is primarily based on voluntary contributions and large-scale collaboration. Peer production is a significant organizational innovation: agents voluntarily self-assign work and successfully coordinate towards the provision of global public goods, in the absence of price signals and without any pre-specified design rule or formal leadership.

From Open Source Software to Wikipedia, peer production involves hundreds of thousands of contributors worldwide. It is an important source of value creation in the most competitive sectors of information and technology, as well as a major source of innovation. Beyond its economic significance, the emergence of peer production also represents an opportunity to shed new lights on a number of longstanding but notably difficult questions in the literature. Given the unconventional nature of many of the work incentives at play in peer production environments, those are particularly well suited for researching the impact of non standard economic preferences on public goods provision, studying their role as work incentives, and assessing their consequences in terms of organizational economics.

The first contribution of this dissertation is a methodological one. Chapter 1 develops and assesses the reliability of a novel online experimentation tool specifically designed to strengthen the internal validity of the decisions elicited over the Internet. Chapter 2 and 3 document the rise of peer production as a new and significant model for organizing production. Exploiting the context of peer production, those chapters leverage the online experimentation tool developed in Chapter 1 and rely on a combination of large-scale online experiments and computational methods (i.e. the systematic extraction of data on subjects’ field behavior) to respectively (i) provide the first comprehensive field test of the theory of the private provision of public goods, (ii) study the importance of social preferences as work motives within real-world productive organizations and (iii) report the first field evidence of endogenous sorting behavior of economic agents within productive teams based on their cooperative types.

Keywords: Field Experiment, Social Preferences, Public Goods, Labor Economics, Peer Production, Wikipedia, Open Source Software, Internet, Methodology

Résumé

Internet est une technologie très attractive pour la mise en place d'expériences. Il permet d'obtenir des échantillons plus grands et plus divers, et donne au chercheur l'opportunité d'extraire d'Internet toute une série de données de terrain qui documentent les décisions et le comportement de ses sujets. Malgré ces caractéristiques attrayantes, le développement du "laboratoire en ligne" en reste à ses balbutiements, principalement du fait des menaces à la validité et des importantes difficultés pratiques liées à l'expérimentation en ligne.

Plus qu'un outil potentiellement puissant pour la mise en place d'expériences, Internet est aussi un terrain de recherche économique très prometteur. Durant les 20 dernières années, sa diffusion a significativement réduit les coûts de communication et augmenté les échanges d'information entre agents économiques. Cette évolution technologique a favorisé l'émergence d'un nouveau modèle de production – la production par les pairs – basée prioritairement sur les contributions volontaires et la collaboration à large échelle. La production par les pairs est une innovation organisationnelle significative: les agents s'auto-assignent des tâches et se coordonnent avec succès vers la production de bien publics globaux, en l'absence de signaux de prix et sans règle de conception préétablie ou leadership formel.

Des logiciels Open Source à Wikipédia, la production par les pairs mobilise des centaines de milliers de contributeurs de par le monde. C'est une source importante de création de valeur dans les secteurs très compétitifs de l'information et de la technologie, ainsi qu'une source majeure d'innovation. Au-delà même de son importance économique, l'émergence de la production par les pairs représente une opportunité d'éclairer un certain nombre de questions anciennes et particulièrement ardues dans la littérature d'un jour nouveau. Compte-tenu de la nature souvent non conventionnelle des incitations au travail dans les environnements de production par les pairs, ceux-ci sont particulièrement adaptés à l'étude de l'impact des préférences économiques non standard sur la production de biens publics, à l'analyse de leur rôle en tant que motivations au travail, ainsi qu'à l'évaluation de leurs conséquences en termes d'économie organisationnelle.

La première contribution de ce travail de thèse est d'ordre méthodologique. Le chapitre 1 développe et évalue la fiabilité d'un nouvel outil d'expérimentation en ligne, construit spécifiquement de manière à renforcer la validité interne des décisions élicitées sur Internet. Les chapitres 2 et 3 documentent l'émergence de la production par les pairs en tant qu'un modèle nouveau et significatif d'organisation de la production. Exploitant le contexte de la production par les pairs, ces chapitres utilisent l'outil d'expérimentation en ligne développé dans le chapitre 1 et s'appuient sur une combinaison d'expériences en ligne à large échelle et de méthodes computationnelles (i.e. l'extraction systématique de données sur le comportement de terrain des sujets) afin de (i) mener le tout premier test de terrain exhaustif de la théorie de la production privée de biens publics, (ii) étudier l'importance des préférences sociales en tant que motivations au travail au sein d'organisations productives réelles et (iii) procéder aux premiers tests de terrain documentant des comportements endogènes d'appariement des agents économiques au sein d'équipes productives en fonction de leur type coopératif.

Mots-clés: Expérience de terrain, Préférences Sociales, Biens Publics, Economie du Travail, Production par les Pairs, Wikipédia, Logiciels Open Source, Internet, Méthodologie