Research Policy 46 (2017) 970–983

Contents lists available at ScienceDirect

Research Policy

journal homepage: www.elsevier.com/locate/respol

Manipulation of explicit reputation in innovation and knowledge exchange communities: The example of referencing in science

Michael A. Zaggl

TUM School of Management, Technische Universität München, Arcisstr. 21, 80333 Munich, Germany

ARTICLE INFO

Article history:
Received 1 October 2015
Received in revised form 26 July 2016
Accepted 28 February 2017
Available online 3 April 2017

Keywords:
Communities
Explicit
Manipulation
Reputation
Reputation-based reward systems
Reciprocity

ABSTRACT

This paper investigates the manipulation of reputation in the context of innovation and knowledge exchange communities. Reputation is crucial for overcoming the free-riding problem and enables community members to be rewarded because their contributions to the common good can be measured. However, the concept of reputation can include the notion of manipulation, which we define as the attempt to change one's reputation without contributing to the community. To investigate the topic of reputation manipulation, we build on the concept of reputation-based reward systems and extend it by distinguishing between *implicit reputation*, which is uncodified, and *explicit reputation*, which is codified and centrally counted. We argue that the possibilities for manipulation differ between these two distinctions. We investigate reputation manipulation empirically in the context of science, which is built on an explicit reputation-based reward system, and we use the received citations as an indicator for reputation. We distinguish two forms of manipulation—unjustified self-citing and unjustified reciprocal citing—and find evidence of both within a bibliometric dataset. This paper contributes to the design of knowledge exchange communities by highlighting the opportunities and challenges arising from explicit reputation-based reward systems, specifically the opportunities for manipulation. It also contributes to the work on misconduct in science.

1. Introduction

Reputation plays a crucial role in innovation and knowledge exchange communities because it can account for the contributions of the communities' members. Reputation has been researched in contexts other than communities, such as organizations and their stakeholders (Lange et al., 2011) and online markets (Bolton et al., 2013; Lanzolla and Frankort, 2016; Moreno and Terwiesch, 2014; Yoganarasimhan, 2013), where a positive reputation increases the likelihood of mutually beneficial transactions. In the context of organizations and their stakeholders, reputation enables the transfer of decision-relevant information to recipients such as potential employees, customers, and suppliers. Similarly, reputation helps to overcome the problem of information asymmetries between buyers and sellers in markets. However, the role of reputation differs in the context of communities of innovation and knowledge exchange such as open source software development (Henkel, 2006; Osterloh and Rota, 2007; von Krogh et al., 2012) and science (i.e., basic research) (Merton, 1988; Stephan, 1996; Stephan, 2010). Here, reputation is based on community members' contributions. Consequently, reputation represents a kind of deposit that enables selective incentives and directs benefits exclusively to the contributors (Milinski et al., 2002; Oliver, 2013; Olson, 1965; von Hippel and von Krogh, 2003). Reputation enables others to reciprocate or the contributor to signal hidden qualities in the expectation of future benefits (Lerner and Tirole, 2002).

The value represented by reputation inevitably incites manipulation attempts (see Charness et al., 2014). We define manipulation as the attempt to change one's reputation without actually contributing to the community. Manipulation is distinct from free riding. Both avoid the cost of contributing, but while free riders are excluded from selective benefits (i.e., benefits exclusive to the contributors, see Olson, 1965), manipulators profit from selective benefits because they pretend they have contributed. The phenomenon of reputation manipulation is almost as old as that of reputation itself. Now, online communities are dictating new conditions, including the possibility to codify reputation information and translate it into an *explicit* representation. This explicit representation of reputation changes the possibilities for manipulation by undermining traditional protection measures such as social punishment for self-praise.

This paper addresses the topic of reputation manipulation in innovation and knowledge exchange communities. It enriches our understanding of such communities (Faraj and Johnson, 2011; Faraj et al., 2011; West and Lakhani, 2008) and the role of explicit reputation. Currently, there is a lack of understanding of the possibilities

http://dx.doi.org/10.1016/j.respol.2017.02.009
0048-7333/

in communities for the manipulation of reputation, resulting in the misuse of selective incentives. Manipulation is a relevant phenomenon as shown, for example, by Sojer et al. (2014) in a recent study on unethical code reuse in open source software development and by Hutter et al. (2015) in innovation contests. Based on the extant literature, we define knowledge exchange and innovation communities as members interacting voluntarily with no common affiliation but with similar goals such as innovation or knowledge accumulation (West and Lakhani, 2008). Community members have a shared language and common rules, and they rely to a large degree on self-governance (Rullani and Haefliger, 2013), and they have the inherited characteristics of collective action systems (see Allen, 1983; Olson, 1965; Ostrom, 2000, 2007).

We chose the scientific system to examine the manipulation of explicit reputation. This system underlies the conditions of collective action and references or citations are a meaningful representation of reputation (Dasgupta and David, 1994; Stephan, 1996). We extract data in the form of a citation network that consists of several thousand publications from the Web of Science. Building on the free-rider hypothesis (Olson, 1965; Hardin, 1968), we expect manipulation to be a significant phenomenon and publications involved in this manipulation to be of inferior quality. Based on these expectations, we develop hypotheses and apply a deductive research design. We find significant evidence for author self-citing and author reciprocal citing in the scientific system. Self-citing is not a new phenomenon (Fowler and Aksnes, 2007; Glänzel et al., 2004; Hyland, 2003), but we extend the concept to include manipulative self-citing. Despite strong interest of the conceptual literature (Phelan, 1999; Posner, 2000), reciprocal citing has rarely been studied empirically at the individual level (i.e., author or publication level), the exceptions being studies using small samples (Paisley, 1990; White et al., 2004). Since self-citation and reciprocal citation are not generally attempts to manipulate reputation, we develop a simple and effective strategy to distinguish alleged manipulation from self- and reciprocal citations that occur without the intention to manipulate (we refer to these as justified self- and reciprocal citations). By controlling for shared content between cited and citing publications, we construct a proxy criterion to distinguish between manipulated and justified citations. Thus, we account for alternative motivations than manipulation. In addition to shedding light on the phenomena of self- and reciprocal citing, our empirical investigation shows that manipulation is associated with inferior quality. This implies that the ability of the scientific system to self-organize (Dasgupta and David, 1994; Martin, 2012) can (partly) correct for manipulation activity.

This paper contributes to the literature in three ways. First, it extends the literature on innovation and knowledge exchange communities (Faraj and Johnson, 2011; Faraj et al., 2011). It provides a new perspective on communities by introducing the dichotomy of implicit versus explicit reputation and by approaching the topic of reputation manipulation. Online communities provide a unique opportunity for implementing governance mechanisms to reduce or even resolve a wide range of cooperation and coordination problems. Exploiting explicit reputation is promising because it increases visibility and efficacy compared to implicit reputation. If reputation is explicit, there is no need for all participants to keep track of all other members' reputations. Instead, reputation is calculated and stored centrally.

We show that manipulation is an issue in such systems, and that governance mechanisms should be designed to minimize the possibilities to manipulate. More generally, this paper provides insights into collective innovation and selective incentives (Oliver, 2013; Olson, 1965; Osterloh and Rota, 2007; von Hippel and von Krogh, 2003). It highlights the downside to selective incentives by relating them to the concept of manipulation. Thus, we challenge the current one-sided, positive characterization of selective incentives by arguing that reputation-based reward systems are inherently accompanied by misleading incentives.

Second, we contribute to the development of a more integrated theory of innovation and knowledge exchange communities. The extant literature includes multiple empirical studies on communities that are isolated (e.g., Franke and Shah, 2003; Lüthje et al., 2005; Jeppesen and Frederiksen, 2006; Antorini et al., 2012; Hienerth et al., 2014). This paper makes a crucial step towards a more holistic and integrated picture of the organizational form of communities by providing a general understanding of reputation and its functioning as one of the main building blocks of communities.

Third, we contribute to the literature on the economics of science (Aghion et al., 2009; Stephan, 1996; Stephan, 2012), specifically the growing body of work on misconduct in science (see Lacetera and Zirulia, 2011; Martin, 2012; Martin, 2013) including the research on citation behavior by bibliometric methods (Fowler and Aksnes, 2007; Hyland, 2003; Kostoff, 1998). Manipulation based on authors' self- and reciprocal citing is more subtle and seemingly less harmful than the currently discussed forms of misconduct such as plagiarism or data fabrication. However, this subtlety might lead to lower concern over these practices and their more widespread use. Thus, the topic of manipulation in the form of self- and reciprocal citing might be underestimated at first glance.

This paper provides some practical implications for the governance of innovation and knowledge exchange communities. First, managers of innovation communities should be aware of the advantages and disadvantages of explicit compared to implicit reputation-based reward systems. Second, community designers and managers should consider the possibilities of reputation manipulation as this could affect the design and governance of existing platforms (e.g., for managing software projects such as GitHub). Third, we recommend policymakers in the scientific systems identify self-citations and distinguish them from other citations.

The remainder of this paper is structured as follows. Section 2 reviews the literature on reputation-based reward systems, elaborates the concept of implicit versus explicit reputation, and develops our hypotheses on manipulation in the form of self- and reciprocal citing in the scientific system. Section 3 describes the dataset and the research design. Section 4 presents the results. Section 5 discusses the findings and provides implications for theory and for practitioners.

2. Theoretical background and hypotheses

This section sets the theoretical background to the paper. Based on the notion of reputation-based reward systems, we distinguish between implicit and explicit reputation to study manipulation in the latter case (Section 2.1). We present reasons for our choice of the scientific system as the context for an empirical investigation (Section 2.2) and develop our hypotheses (Section 2.3).

2.1. Implicit and explicit reputation-based reward systems

Reputation-based reward systems are elegant solutions for public good and collective action problems. Reputation accounts for contributions and enables mechanisms that compensate individuals for innovating and revealing their private knowledge (Dasgupta and David, 1994; Stephan, 1996; von Hippel and von Krogh, 2003). Reputation-based reward systems can differ in terms of their representation of reputation. We introduce a distinction between *implicit* and *explicit* reputation.

We introduce the term *implicit reputation-based reward system* to refer to systems that rely on reputation but lack central coordination of reputation-related information with the result that every individual accounts for the reputation of the other individuals on his/her own. Implicit reputation-based reward systems can be seen as the original form for governing reputation under collective action and public good conditions. They occur mainly in small, close-knit communities (see Ostrom, 2000, 2007) where everyone can observe everyone else. The individual tracking of others' reputations requires extensive cognitive capacities (see Nowak and Sigmund, 2005). These capacities limit the number of members participating in the community. Further, fluctuations in community members can disturb the efficacy of reputation.

The contrasting concept is an *explicit reputation-based reward system* through which explicitness of reputation refers to the condition that the information represented by reputation is codified and consolidated. Instead of individual and decentralized tracking of all participants' reputations by all other participants, a central institution calculates and keeps track of the reputations. Explicit reputation-based reward systems seem to have several advantages over implicit reputation-based reward systems. Experiments show that people do not use reputation optimally because of cognitive limitations. Instead, they apply simpler, but less advantageous, reputation-based strategies (Milinski et al., 2001). It is likely that explicit reputation-based reward systems have the potential to overcome these problems. They reduce the cognitive capacity requirements of individuals. Also, explicit representation of reputation reduces uncertainty about the value of contributions to the common good, and a system based on explicit reputation is likely to be more robust to the interference of noise. In contrast, individual calculations of others' reputation in implicit reputation-based reward systems are more likely to be inaccurate and contradictory. Thus, explicit reputation reduces possible disagreements over status differences, and this is desirable because such disagreements are detrimental to contribution behavior (see Kilduff et al., 2016).

Distinguishing explicit from implicit reputation triggers a fundamental discussion about the incentives in innovation and knowledge exchange communities (see Faraj and Johnson, 2011; Faraj et al., 2011; Boudreau and Lakhani, 2015). Communities are a crucial source of innovation, and examples are open source software development (Dahlander and Magnusson, 2005; Henkel, 2006; Hertel et al., 2003; Stam, 2009; von Krogh et al., 2003) and idea crowdsourcing (Piezunka and Dahlander, 2015; Schemmann et al., 2016). They enable explicit reputation and this raises the question about the optimal design of reputation-based reward systems, especially when considering manipulation. In implicit reputation-based reward systems, other participants react negatively to self-praise and related forms of manipulation, and thus correct for such manipulation attempts. These social control mechanisms have evolved in implicit reputation-based reward systems but they lose their effectiveness in explicit reputation-based reward systems.

Despite the crucial role of explicit reputation-based reward systems for solving collective action and public good problems in communities, they are rarely researched in this context. Most research is in the context of online markets for services, goods, and labor, e.g., Amazon, eBay, and Upwork (Bolton et al., 2013; Dellarocas, 2006; Moreno and Terwiesch, 2014; Yoganarasimhan, 2013). However, the insights from that research cannot be generalized to communities because the role of reputation differs fundamentally. While reputation helps to overcome information asymmetries between buyers and sellers in markets, it represents information on the contributions to the common good in the context of communities. Science is virtually the only domain in the context of collective action and community-based innovation in which reputation has been researched in depth (Dasgupta and David, 1994; Stephan, 1996).

2.2. Science as the research context

For several reasons, we chose to investigate the question of reputation manipulation in the scientific system. The scientific system shares the basic characteristics of innovation and knowledge exchange communities (Fleming and Waguespack, 2007) and it relies on an explicit reputation-based reward system. Scientists pursue common goals and interact with each other for that reason (cf., West and Lakhani, 2008). They produce knowledge and relinquish control over this knowledge (Dasgupta and David, 1994; Merton, 1973; Stephan, 1996). To provide incentives and to coordinate, scientists, as well as the members of innovation and knowledge exchange communities, rely on self-governance based on shared rules and norms (cf., Fauchart and von Hippel, 2008; Rullani and Haefliger, 2013) in which reputation is a crucial element (e.g., Powell, 1990).

In many innovation and knowledge exchange communities, reputation is reflected by the voting of other members. This reputation usually is codified and easily available to interested individuals; thus, reputation is explicit. Based on their reputation, community members can gain benefits from inside and outside the community (e.g., Lerner and Tirole, 2002; Wasko and Faraj, 2005). Analogously, in the scientific system, citations resemble voting and provide a similar measure of reputation. They represent the reuse of the published knowledge and reflect the value of the cited contributions (Merton, 1988; Stephan, 1996). Like community members, scientists are rewarded for their contributions by being granted access to diverse resources such as research funds, prizes, honorary degrees, invitations, and tenure. A direct connection between citations, reputation, and reciprocal behavior from the community is the example of the Nobel Prize, the most important award in physics, medicine, chemistry, and economics. A scientist's citation count is a crucial predictor of the next laureate (Garfield, 1986; Garfield and Welljams-Dorof, 1992). Reputation in science is also explicit. Databases such as the Web of Science, Scopus, and Google Scholar codify publications and citations and aggregate the information.

The extant research on the scientific system provides a basis for investigating the topic of reputation manipulation. This research covers different forms of misconduct, such as data fabrication and plagiarism, including self-plagiarism (see Enders and Hoover, 2004; Lacetera and Zirulia, 2011; Martin, 2013). Relatedly, the extant bibliometric research addresses the phenomenon of self-citation (Aksnes, 2003; Fowler and Aksnes, 2007; Glänzel and Thijs, 2004; Glänzel et al., 2004; Hyland, 2003; Phelan, 1999; Snyder and Bonzi, 1998; van Raan, 2008; for a review, see Glänzel et al., 2006). This research finds that self-citing is common, but the intensity varies across scientific disciplines (Aksnes, 2003; Snyder and Bonzi, 1998). However, self-citing does not necessarily fulfill our definition of manipulation. It is essential to know the motives of self-citing authors to infer such a connection because not all self-citations are intended merely to push up citation counts. Authors often cite their own prior work in order to reference its contents. Sometimes it would be unethical for them not to cite earlier works. Based on surveys, extant research has investigated the motivations for self-citing (e.g., Bonzi and Snyder, 1991). However, the results of these studies are inconclusive and methodologically difficult to interpret because the authors had no reason to report unethical motives such as manipulation.

The extant bibliometric literature also discusses reciprocal citing. However, this part is much smaller than the work on self-citing and is almost exclusively conceptual (Kostoff, 1998; Phelan, 1999; Posner, 2000), except for a few empirical studies that point to the occurrence of reciprocal citing (Paisley, 1990; White et al., 2004). These studies use very small samples (less than 20 authors). Similar to the case of self-citing, authors' motivations for reciprocal citing are difficult to analyze. It is therefore, not trivial to distinguish manipulative from justified reciprocal citations.

Altogether, the scientific system seems suitable for approaching the topic of reputation manipulation. The extant bibliometric literature hints to self- and reciprocal citing as possible candidates for identifying manipulation. The literature shows empirically that these citation practices occur. However, it does not associate the intention of manipulation with self- and reciprocal citing.

The common problem in measuring manipulation is the lack of information on the motives of alleged manipulators. It is difficult to distinguish justified motives (referencing one's own prior content) from manipulative reasons (increasing one's citation count). Interestingly, authors' being aware of this difficulty may intensify their manipulative behavior. An author can claim that a citation was intended to reference prior knowledge contained in his/her own previous work or the work of his/her friends. This creates a kind of immunity, which, in turn, may increase the temptation to manipulate.

To approach reputation manipulation in science, it is necessary to differentiate between authors' motives for citing. We approximate this differentiation by isolating publication content and using it as a rough indicator: A self- or reciprocal citation without content overlap of cited and citing publications is interpreted as being more likely aimed only at increasing the cited author's reputation, which we refer to as unjustified self- or reciprocal citation (manipulation). Conversely, a self- or reciprocal citation with content overlap of cited and citing publications is interpreted as being more likely to be motivated by the intention to refer to prior knowledge, thus we denote it as justified self- or reciprocal citation (no manipulation). Although the use of content as an indicator seems to be insufficient to infer the motivations of particular authors, it allows us to draw conclusions at the aggregate level. Simply put, citations between publications without content overlap are *more likely* to be cases of manipulation compared to citations connecting publications with shared content.

2.3. Hypotheses development

In what follows, we develop hypotheses on reputation manipulation in the form of self- and reciprocal citing in the context of the scientific system (Section 2.3.1). We also develop hypotheses on the quality of publications receiving self- and reciprocal citations (Section 2.3.2).

2.3.1. Identifying manipulation of reputation

Manipulation of a reputation-based reward system refers to capturing value, i.e., reputation and the associated benefits, without contributing to the public good, and thus not creating value. Manipulation is similar but not identical to traditional free riding, which refers to not contributing, while manipulation consists of pretending to contribute but actually withholding contributions. Reputation manipulation disturbs selective incentives. Contributors have to share their benefits with the manipulators. Exceeding a certain level, manipulation can lead to a collapse of the entire reputation-based reward system.

We tackle manipulation in explicit reputation-based reward systems from the perspectives of the free-rider hypothesis (Olson, 1965; Hardin, 1968). This theory predicts manipulation behavior in collective action. It explains the under-maintenance of public goods and the overexploitation of commons (Hardin, 1968). People are unwilling to contribute to a public good or commons without further incentives (Olson, 1965). These incentives must be directly linked to contributions; in other words, they are selective (or exclusive) for contributors and only they can claim the benefits. Thus, the contributors are motivated by the expectation of returns from these selective benefits. In the absence of expected returns directly linked to contributions, the individual saves on contribution costs. In the context of reputation manipulation in communities, the free-rider hypothesis suggests that participants will fake their contributions whenever they can. The manipulator can profit from benefits exclusive for participants with reputation but without having to bear the cost. From this perspective, the possibility to manipulate creates another public good/collective action problem which people are willing to exploit.

In science, the possibility to cite one's prior work represents an opportunity to manipulate. Self-citing is a common phenomenon in the scientific system (Aksnes, 2003; Bonzi and Snyder, 1991; Fowler and Aksnes, 2007). However, as mentioned, self-citation counts per se are insufficient to identify manipulation since they include justified self-citations. Authors following their own research agendas will cite their prior publications on the basis of common content. This circumstance makes controlling for content necessary. If authorship beyond content can predict citations, then self-citing is aimed at shifting one's reputation. This leads to the following hypothesis:

Hypothesis 1a (H1a). Shared authorship, beyond shared content, of two publications increases the probability of a citation between them.

Similar to self-citing, we expect that authors engage in reciprocal citing. Authors can increase their citations by citing each other's publications reciprocally. As mentioned, it has been empirically shown that authors engage in this practice (Paisley, 1990; White et al., 2004) and various explanations for this phenomenon have been proposed (Kostoff, 1998; Phelan, 1999; Posner, 2000). However, these studies do not address reciprocal citing motivated by manipulation.

Extant theories in biology, sociology, and economics characterize reciprocity as a strategy that is mutually beneficial for its participants (Axelrod, 1984; Fehr and Gächter, 2000; Gouldner, 1960; Trivers, 1971). Receiving a benefit from another party obliges the receiver to return the favor (Gouldner, 1960). In the context of the scientific system, this means that an author who has been cited will feel inclined to return the favor. This benefits all the authors involved. Thus, from the perspective of the free-rider hypothesis, we predict that authors apply the principle of reciprocity to strategically inflate their citation counts.

As in the case of self-citation, controlling for content is necessary in an examination of reciprocal citing behavior. Authors working in the same field will automatically, and without the intention to manipulate, engage in reciprocal citation relations. This type of citation is described as justified reciprocal citation, not manipulation, and must be treated accordingly.

Hypothesis 1b (H1b). Authors who have been cited by some other author reciprocate the citation beyond referring to shared content.

2.3.2. Quality characteristics of manipulation

We develop two hypotheses about the quality effects by manipulation. Our reasoning is based on the abilities of the scientific system to govern itself (Dasgupta and David, 1994). We expect that these abilities to (partly) counterbalance the effects of a manipulated reputation, i.e., a reputation from self- and reciprocal citing. Manipulation can negatively influence the quality of a publication, and the scientific system is sensitive to such quality shifts. The basic measure of publication quality is the citation count (Merton, 1988; Stephan, 1996). Thus, we expect that manipulation is likely to be correlated with publication quality measured as the number of citations received. We propose two reasons for the relation between manipulation and poor quality.

First, the motivations of authors who are inclined to manipulate are likely extrinsic rather than intrinsic (see Deci and Ryan, 1985; Gagné and Deci, 2005), and extrinsic motivation might lead to lower publication quality. Extrinsically motivated authors tend to avoid costs and thus minimize efforts when developing their publications. The reduced effort reduces quality. In contrast, an intrinsically motivated author will be more diligent because of immediate returns from working, resulting in quality improvements for the publication. Thus, manipulation is related to lower quality.

Second, manipulation is more beneficial for poor performers than high performers. The gains from manipulation are higher for authors whose publications have low citation counts. Thus, we expect that inferior quality characterizes manipulation. Alternatively, it could be argued that high-performing authors have a stronger tendency to manipulate. They have more publications and more opportunities to self-cite. They are also socially better connected in the scientific community and therefore have more possibilities for reciprocal citing. However, it is likely that authors with many publications limit their self-citations by referencing only to those publications with matching content and thus they avoid manipulations.

In addition, the positive feedback mechanism in the scientific system that is known as the Matthew effect can help to explain why manipulation leads to fewer citations. Receiving citations increases awareness of the cited publication and this attracts more citations (Merton, 1968, 1988). Many authors, however, scrutinize the publications they become aware of and reject citing low-quality publications. Thus, we expect that manipulated citations do not have the same potential for causing further citations of a cited publication. Their influence on future citations should be weaker or even negative compared to the effect of regular citations.

Thus, we expect that publications receiving self-citations motivated by manipulation will be less often cited compared to publications receiving self-citations not motivated by manipulation or publications receiving citations from other authors. Again, controlling for content helps to distinguish justified self-citations (non-manipulative) from unjustified self-citations (manipulative).

Hypothesis 2a (H2a). Receiving unjustified self-citations, relative to receiving other types of citations, is associated with a lower number of received citations.

Similarly, it is expected that reciprocally cited publications receive a smaller total number of citations. Their quality is inferior compared to publications with citations that are not aimed at manipulation. Here again, controlling for content allows a distinction to be made between justified and unjustified reciprocal citations.

Hypothesis 2b (H2b). Receiving unjustified reciprocal citations, relative to other types of citations, is associated with a lower number of citations received.

3. Dataset and method

3.1. Dataset

We collect publication data, including citation relationships, from Thomson Reuters' Web of Science. The data collection strategy specifically aims to capture chains of citation relationships. We use a similar procedure as Belenzon (2012). The procedure results in a citation tree consisting of a root publication and its immediate and successive citing publications (i.e., the publications citing the root, and then the publications citing these publications, and so on). Such a citation tree captures multiple citation relationships and long citation chains. It also minimizes the number of decisions required: We selected only the root publication, and so the rest of the network is determined.

Table 1
Overview of the dataset.

Publications	61,978
Journals	4,184
Authors per publication	5.32
Keywords per publication	12.06
Words in abstract per publication	116.75
Publications (only type *article*)	61,335
Citing publications	61,977
Cited publications	25,499
Citation chains of length 2	110,479
Citations with overlapping authors (self-citations)	23,596
Cited A-publications in chains of length 2 (only type *article*)	24,912
Citations received by A-publications in chains of length 2 (only type *article*)	107,160
Mean total number of citations received of publication A[a]	4.26
Mean total number of citations received of publication A (only publications that received at least one self-citation)	6.02
Mean journal-weighted number of citations received[b] of publication A (only publications that received at least one self-citation)	1.55
Citation chains of length 3	142,188
Citation chains of length 3 without self-citations	93,298
Citation chains with overlapping authors (between first and last publication; that is, reciprocal citations)	5,746
Citation chains of length 3 (only type *article*)	135,054
Cited B-publications in chains of length 3 (only type *article*)	33,314
Citations received by B-publications in chains of length 3 (only type *article*)	88,840
Mean total number of citations received of publication B	4.10
Mean total number of citations received of publication B (only publications that received at least one reciprocal citation)	5.06
Mean journal-weighted number of citations received of publication B (only publications that received at least one reciprocal citation)	0.63

[a] Note that these publications are cited at least once. Otherwise, they could not be a cited publication in a citation chain.
[b] For a definition of the journal-weighted citations received, see Section 3.2.1.

To choose the root publication, we used a 10-year time span, which we considered sufficient to collect a citation tree of suitable size. We judged that data from a scientific field with frequent publications would be the most meaningful and opted for the field of biochemistry. We randomly selected a publication from the issue of *Journal of Biochemistry* that has been published exactly 10 years earlier.[1] Collecting the citation network data took several months and was conducted over the fourth quarter of 2012 and the first quarter of 2013. Because of the data structure and the multiple types of relationships (e.g., relationships between different publications, and relationships between publications and authors), we transferred the data into a graph database.[2]

Table 1 provides an overview of the dataset. It shows that the dataset contains a large number of publications (61,978) and covers a wide range of journals (4,184). The major journals are *PLoS One* (6.1% of the publications), *Molecular Ecology* (1.8%), and *PNAS* (1.4%). The data include author names, keywords, and abstracts. Typical for the natural sciences is a large number of authors per publications (5.32). The Web of Science distinguishes among types of publications (*reviews*, *conference papers*, *editorials*, etc.).

[1] This is Nong V.H., Arahira M., Phan V.C., Kim C.S., Zhang D., Udaka K., Fukazawa C., 2002. Molecular cloning and characterization of a group II chaperonin delta-subunit from soybean, *Journal of Biochemistry*, 132(2):291–300.
[2] A graph database stores information in the form of graphs in contrast to relational databases, which use tables. We used the Neo4j database management software (neo4j.com).

For testing H1a and H1b, we use citation chains of length 2 (publication A is cited by publication B) and length 3 (publication A is cited by publication B and publication B is cited by publication C, but publication C does not cite publication A). Citation chains of length 2 (n = 110,479) allow us to examine self-citing. Citation chains of length 3 allow us to analysis reciprocal citing. To analyze reciprocal citing, we exclude all cases of self-citing and the sample size is n = 93,298.

For testing H2a and H2b, we consider only the type of *article*, which represents the majority of all publications. Taking publications as unit of analysis requires controlling for the type of publication because *reviews* usually receive more citations, whereas *conference papers* and other forms tend to receive fewer citations. For analyzing self-citing (H2a), we focus on the first publications (A-publications) in citation chains of length 2 (n = 24,912) and the citations they received (n = 107,160). The analysis for reciprocal citing (H2b) focuses on the B-publications in citation chains of length 3 (n = 33,314) and the citations they received (n = 88,840).

3.2. Method

Our empirical strategy to test H1a and H1b is based on identifying the drivers of citation relationships. Simply put, is authorship, beyond the publication's content, a citation driver? Therefore, we mix citation chains of the length 2 (length 3) with non-existing, but theoretically possible, citation chains. We apply logistic regressions to measure the predictors of the citation relations.

To measure the quality of publications (H2a and H2b), we take publications as unit of analysis and measure the types of citations early in the lifetime of a publication that lead to more or fewer citations received at a later date. Simply put, do self-citations or reciprocal citations lead to more or fewer citations received? Therefore, we focus on publications and the number and type of their citations. For both strategies, we operationalize the following variables.

3.2.1. Dependent variables

For H1a and H1b, the dependent variable is the binary information whether a citation relation between publications exists or not. When testing the hypotheses, we will add data on non-existing citation relations.

H2a and H2b focus on the quality of the publications that received self- or reciprocal citations. We operationalize publication quality as the *total number of citations received* and the *journal-weighted number of citations received*. The total number of citations received is the number of all citations received by the focal publication (c_i, where i is the focal publication). The journal-weighted number of citations received is the focal publication's total number of citations received minus the mean number of citations received by all the publications published in the same journal as the focal publication ($c_i - \frac{\sum_{j=1}^{|J|} c_j}{|J|}$, where J is the set of all publications printed in i's journal). This represents a publication's expected number of citations received contingent to its journal. We further specified both variables by distinguishing the time distance between the focal publication (i) and the citing publications, and whether there is an overlap among authors and among content.

3.2.2. Independent variables: author similarity

We define authorship similarity of a publication pair as the main independent variable. Similarity (or overlap) is operationalized in two ways. First, we calculate a binary measure distinguishing whether both publications have any authors in common or whether there is no author overlap (Carley et al., 2013; Snyder and Bonzi, 1998). The second and a more granular measurement provides the Jaccard similarity coefficient (see Schubert et al., 2006). The Jaccard coefficient can be calculated for any pair of publications as the ratio of the number of overlapping authors to the total number of unique authors (i.e., if a and b are the sets of authors of the publications A and B, then $jaccard_{authors}(A, B) = \frac{|a \cap b|}{|a \cup b|}$). Note that for the extremes $\{0,1\}$ Jaccard similarity coefficient and binary match are identical.

3.2.3. Control variables: content similarity

To measure content similarity between two publications, we use three indicators. First, we compare the keywords. For each publication pair, we calculate the binary overlap similarly to our measure of author similarity. If there is overlap in any of the keywords, the binary measure indicates a match; otherwise, there is no match. We also calculate the Jaccard similarity coefficient for the keywords. Similar to the Jaccard similarity of authors, we divide the number of identical keywords by the total number of unique keywords of the two publications (i.e., if a and b are the sets of keywords of the publications A and B, then $jaccard_{keywords}(A, B) = \frac{|a \cap b|}{|a \cup b|}$).

Second, we use the journal as an indicator of content overlap. Therefore, we calculate a binary variable whether two publications were published in the same journal or not.

Third, we use the publication abstracts as an indicator for content similarity. To compare the abstracts of two publications, we followed the procedure by Piezunka and Dahlander (2015). We queried our graph database and then imported the results into the statistical software R. Using the packages *quanteda* and *lsa*, we transformed the text of the abstracts into lowercase capitals. Words without specific content (so called stopwords, e.g., a, an, the, no) and general terms (e.g., abstract, specific, we) were removed as were numbers and punctuation. We then converted the remaining words into their stem forms. For each publication, we calculated a vector of each word's number of occurrence. We allocated higher weights to rarely used words. We assume that two abstracts are more similar if both use a rarely occurring word compared to if they share a common word. Therefore, each word was weighted according to: *log(number of word occurrence in the focal abstract + 1) · log(1 + (number of abstracts/number of abstracts using the word))*. Eventually, the cosine similarities between the vectors were calculated, representing a measure of the similarity between two publications (cf., Piezunka and Dahlander, 2015).

4. Results

4.1. Identifying manipulation of reputation

4.1.1. Identifying self-citing (H1a)

We examine self-citations by estimating the probability of a citation relationship between two publications based on the predictors of similarity in authorship and content. The unit of analysis is the citation chain of length 2 (n = 110,479, cf., Table 1). In addition to these citation chains (positive cases), we consider data on non-existing but theoretically possible citation relations (negative cases). In principle, a citation relation is possible in any publication pair as long as the citing publication is younger than the cited publication. However, considering all possible citation relations as negative cases is impossible for reasons of computational limitations (61,978 publications, cf., Table 1, lead to $61978 \cdot (61978 - 1)/2 \approx 1.9 \cdot 10^9$ possible citation relations). Thus, we reduce the number of negative cases by selecting a random subset. We decide to use a subset that is three times as large as the number of positive cases. We merge positive and negative cases, and removed cases without keywords or abstracts, which results in $n = 110{,}479 + 3 \cdot 110{,}479 - 2{,}429 = 439{,}487$ cases, each representing a pair of publications with binary information whether a citation relation between them exists or not and the information on the similarities of authors, keywords, journals, and abstracts.[3]

Table 2 presents the results of the logistic regression analyses, which support H1a. Models 1 and 2 use the binary measures for the similarity of authors and keywords (see Section 3.2). Model 1 uses only content-based predictors (keywords, abstracts, and journals). Model 2 adds author similarity and shows that it is a strong indicator for the existence of a citation relation beyond the content measures. Author similarity is stronger than keywords and journal but weaker than the similarity between abstracts.

Regression models 3 and 4 use the more granular Jaccard similarity coefficient for authors and keywords. Compared to the binary measures of models 1 and 2, the Jaccard measure shows a stronger impact of keywords in model 3 and a stronger impact of keywords and authors in model 4.[4] Model 4 shows that shared authorship is the strongest predictor, stronger even than similarity between abstracts.

4.1.2. Identifying reciprocal citing (H1b)

To examine reciprocal citing, we use citation chains of length 3 and test whether the last publication has an author in common with the first publication. We also consider content as a citation driver. The following describes our approach. If publication A is cited by publication B, and publication B is cited by publication C, and there is no citation of publication A by publication C, we define this as a citation chain of length 3 (cf., Section 3.1). We want to infer possible reciprocal behavior by the authors of publication C when citing publication B. Reciprocal citing is the case only if authors of publications A and C are identical. These authors (or one of them) reciprocated the citation by the authors of publication B. This is not the only way to conceptualize reciprocity. For example, the authors of publication C, rather than citing publication B, can reciprocate by citing another publication authored by publication B's authors. In our analysis, we are agonistic about alternative reciprocation possibilities, and only focus on the evidence identified by the described approach.

We control for content similarity (using keywords, journals, and abstracts) between publications B and C as an alternative reason (non-manipulation) for the authors of publication C to cite publication B. To avoid interference from self-citing, we excluded all citation chains with an author overlap between publications A and B or B and C.

We consider all citation chains of length 3 from the data matching the criteria defined and thus obtain 93,298 citation chains (positive cases) (cf., Table 1). Similar to the procedure used to test H1a, we add data on non-existing publication relations (negative cases). We do so by adding citation chains where publication A is cited by publication B, but publication C does not cite publication B. All other criteria defined for the positive cases apply to the negative cases. We take a random subset from all possible negative cases that is three times as large as the number of positive cases and remove cases without keywords or abstract information ($n = 93{,}298 + 3 \cdot 93{,}298 - 4158 = 369{,}034$).[5]

Table 3 shows the results of the logistic regression analyses testing reciprocal citing. The results support H1b. Model 1 uses journal, abstract, and the binary similarity measure for the keywords as predictors. As expected, content similarity predicts the citations between publications B and C. Model 2 includes binary author similarity (of publications A and C) as a predictor. Author similarity is a significant but relatively weak predictor compared to the measures of content similarity, especially the abstract. Models 3 and 4 use the Jaccard similarity coefficient instead of the binary measurement for keywords and authors. Here, the effect size of keywords and authors is larger. In model 4, author similarity is a significant predictor but weaker than similarity in keywords and abstract.

4.2. Effects on quality

4.2.1. Self-citing and quality (H2a)

We hypothesized that publications that benefit from reputation manipulation are associated with inferior quality. Specifically, H2a states that manipulatively self-cited publications are less often cited than publications that received other types of citations. To test this hypothesis, we distinguish between *justified* and *unjustified self-citations*. Justified self-citations link two publications with common authors *and* have some shared content (here operationalized as overlap in keywords).[6] Unjustified self-citations connect two publications with no shared content but common authors. We compare the impact of the type of citation (justified and unjustified self-citations, and citations from other authors) on the total number of citations received. This approach relies on the stylized fact of positive feedback from citations: receiving citations attracts additional citations (Merton, 1968, 1988). H2a is supported if unjustified self-citations attract fewer additional citations than the other types of citations.

We consider all cited publications, in other words, all cited A-publications in citation chains of length 2 of the type *article* ($n = 24{,}912$). These publications received 7,140 unjustified self-citations, 15,740 justified self-citations, and 84,280 citations from publications of the type *article* (in sum 107,160 citations, cf., Table 1).

First, we compare the shares of justified and unjustified self-citations relative to the total number of citations received (and the journal-weighted number of citations received). Fig. 1 shows that both types of self-citation are more prevalent in rarely cited publications than in publications that are often cited. In other words, their share declines with an increasing number of citations received.[7] This decline in self-citations is not necessarily an indication of a poorer quality. The higher the number of total citations received, the lower the chance that the citations to a publication come mainly from the same author. However, differentiating between justified and unjustified self-citations allows us to infer their relation to quality relative to each other. Both types of citations are subject to the same decline in the chance of being cited by the same author. Fig. 1 suggests that unjustified self-citations are more likely linked to rarely cited publications—they decline more rapidly with increasing quality of the focal publication. The often cited publications have a higher share of justified self-citations. This difference is shown in both quality measures: the total number of citations received and the journal-weighted number of citations received. The difference is supported statistically: publications with a higher share of unjustified self-citations compared to justified self-citations receive fewer citations in total than the publications with a higher share of justified self-citations compared to unjustified self-citations ($M_{\text{moreunjustifiedthanjustifiedself-citations}} = 4.58$, $M_{\text{morejustifiedthanunjustifiedself-citations}} = 6.05$; $t = 3.533$, $p < .001$). The journal-weighted measure of received citations sup-

[3] Robustness checks showed that the results are stable over large variations of the dataset. To test robustness, we changed dataset size and selection of the theoretically possible but non-existing citation relations.

[4] In model 4, the Jaccard similarity coefficient leads to a quasi-perfect separation. To overcome this problem, we applied Firth's bias reduction to penalize the maximum likelihood values (see Firth, 1993; Heinze and Ploner, 2003).

[5] Again, we tested the robustness of our results by changing the selection of the negative cases and found that the results are stable.

[6] See Appendix A for a robustness test that shows that the results do not depend on the choice of keywords as the only indicator to distinguish between justified and unjustified self-citations.

[7] This finding is consistent with the results in Aksnes (2003).

Table 2
Logistic prediction of citation relationship in citation chains of length 2 (self-citing).

	Model 1		Model 2		Model 3		Model 4	
	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*z*-value
$Authors_{binary}$			7.75*** (0.20)	38.2				
$Authors_{Jaccard}$							198.56*** (6.74)	29.5
$Keywords_{binary}$	2.68*** (0.02)	146.7	2.55*** (0.02)	134.5				
$Keywords_{Jaccard}$					47.47*** (0.35)	135.2	45.50*** (0.36)	125.4
Journal	4.35*** (0.03)	139.7	4.20*** (0.03)	129.2	4.38*** (0.03)	139.3	4.20*** (0.03)	128.3
Abstract	44.82*** (0.19)	237.1	45.40*** (0.20)	228.1	46.23*** (0.19)	239.2	46.61*** (0.20)	231.2
Constant	−4.06*** (0.01)	−355.4	−4.20*** (0.01)	−346.9	−4.10*** (0.01)	−352.7	−4.22*** (0.01)	−345.5
Log-likelihood	−88,082.4		−81,011.6		−87,074.7		−80,774.5	
n	439,487		439,487		439,487		439,487	
Nagelkerke's R^2	.763		.786		.766		.787	

*** $p < .001$.

Table 3
Logistic prediction of citation relationship in citation chains of length 3 (reciprocal citing).

	Model 1		Model 2		Model 3		Model 4	
	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*z*-value
$Authors(A,C)_{binary}$			4.16*** (0.22)	18.7				
$Authors\ (A,C)_{Jaccard}$							31.79*** (2.65)	12.0
$Keywords\ (B,C)_{binary}$	2.14*** (0.02)	126.6	2.14*** (0.02)	126.1				
$Keywords\ (B,C)_{Jaccard}$					37.37*** (0.31)	120.2	37.28*** (0.31)	119.8
Journal (B,C)	4.38*** (0.05)	80.0	4.37*** (0.05)	79.6	4.39*** (0.06)	79.9	4.38*** (0.06)	79.5
Abstract (B,C)	39.68*** (0.16)	241.0	39.67*** (0.16)	240.4	40.75*** (0.17)	244.2	40.73*** (0.17)	243.8
Constant	−3.67*** (0.01)	−366.8	−3.68*** (0.01)	−336.5	−3.69*** (0.01)	−335.0	−3.70*** (0.01)	−334.8
Log-likelihood	−87,114.1		−86,822.8		−86,695.7		−86,468.3	
n	369,034		369,034		369,034		369,034	
Nagelkerke's R^2	.712		.713		.714		.715	

*** $p < .001$.

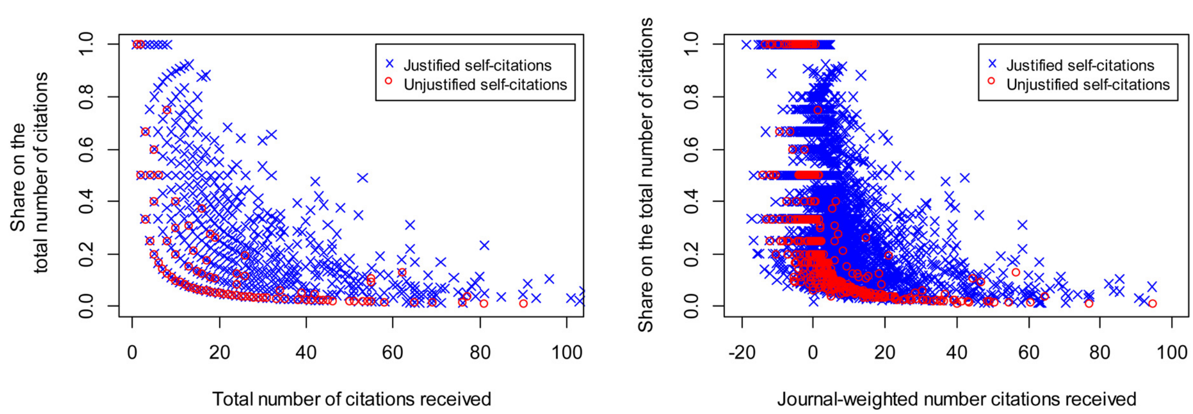

Fig. 1. Share of justified and unjustified self-citations in relation to citations received; left: total number of citations received; right: journal-weighted number of citations received.

ports this result: $M_{\text{moreunjustifiedthanjustifiedself-citations}} = -0.95$, $M_{\text{morejustifiedthanunjustifiedself-citations}} = 1.60$; $t = 6.270$, $p < .001$.

Second, we test the effect of different types of citations in a regression analysis. We use the types of citation (justified and unjustified self-citations and citations from other authors) received by a publication in the same year or one year after its publication to predict citations received after that period. We compute the logarithms of the predictor variables to manage dispersion and specify two negative binomial regressions and one ordinary least square (OLS) regression (see Table 4, models 1–3).[8] Model 1 predicts the total number of citations received in the second year after publication and later. This model shows a positive influence of all the types of citations. Unjustified self-citations have the weakest impact, and citations from other authors are the strongest. Model 2 predicts the number of citations received from other authors only (i.e., we excluded all self-citations from the dependent variable) in the second year after publication and later. In this model, unjustified self-citations have a negative impact. Justified self-citations have no significant influence. Citations by other authors have a positive influence. Model 3 uses the journal-weighted number of citations received in the second year after publication and later as the dependent variable. This variable is a continuous measure, thus we calculated an OLS regression. The model shows that unjustified self-citations have a negative impact. Justified self-citations and citations from other authors are positively related to the journal-weighted number of citations received.

Overall, we find that citations by other authors increase citation counts. This is consistent with the theory and earlier research (Merton, 1968, 1988). The result of model 2, which measures only the impact on citations from other authors, is interesting, as it shows that neither type of self-citation has the potential to attract citations from other authors. This finding is inconsistent with earlier research by Fowler and Aksnes (2007), who suggest a positive effect of self-citations not only on the total number of citations received but also on the number of citations received from other authors. The reason for the difference might be that Fowler and Aksnes (2007) take the authors as a unit of analysis and our results are based on the publication level. If that is the reason, it could be argued that although self-citations attract fewer citations to a given publication (compared to citations from other authors), the author can still profit from self-citing because it spreads his/her name and his/her other publications receive citations.

Most importantly, we observe a negative influence of unjustified self-citations on the citations received from other authors (model 2) and on the journal-weighted number of citations received (model 3). In predicting the total number of citations (model 1), unjustified self-citations have a positive but comparatively weak influence. Thus, H2a is supported.

4.2.2. Reciprocal citing and quality (H2b)

We test the quality of the publications that received reciprocal citations in a similar way as we tested H2a. First, we distinguish between justified and unjustified reciprocal citations. A reciprocal citation occurs when an author cites a publication that cited one of his/her earlier publications. That is, publication A is cited by publication B, publication B is cited by publication C, and publications A and C have a common author but publication C does not cite publication A. In addition, publications A and B and publications B and C do not have common authors (cf., Section 4.1.2). We distinguish between justified and unjustified reciprocal citations. Unlike unjustified reciprocal citations, justified reciprocal citations have some common content (overlap in keywords of publications B and C). We examine the degree to which the different types of citations (justified and unjustified self-citations and citations from other authors) influence the total number of citations received. If unjustified reciprocal citations attract fewer further citations than the other types of citation, we can support H2b.

We use the same data that we used to test H1b, except that we exclude all publications other than the type *article* and focus on publications (specifically publication B) as unit of analysis. The sample consists of $n = 33{,}314$ potentially reciprocal citing publications. We find 373 unjustified reciprocal citations, 841 justified reciprocal citations, and 87,626 non-reciprocal citations (i.e., citations from other authors) in total 88,840 citations of type *article* (cf., Table 1).

Similarly to the test of H2a, we visualize the shares of justified and unjustified reciprocal citations in relation to the total number of citations received and the journal-weighted number of citations received (see Fig. 2). We find that the share of justified and unjustified reciprocal citations declines for publications that received a higher total number of citations. Differentiating between justified and unjustified reciprocal citations does not clearly show that the latter's decline is stronger for the total number of citations received or for journal-weighted number of citations received. We find no statistical difference between the impact of justified and unjustified reciprocal citations (total number of citations received: $M_{\text{more unjustified than justified reciprocal citations}} = 4.60$, $M_{\text{more justified than unjustified reciprocal citations}} = 5.25$; $t = 0.910$, $p = .363$; the journal-weighted number of citations received: M more unjustified than justified reciprocal citations = –0.14, $M_{\text{more justified than unjustified reciprocal citations}} = 0.96$; $t = 1.637$, $p = .103$).

Similar to the method used for H1b, we test the effect of different types of citations with regression analyses. We use the three types of citations (justified and unjustified reciprocal citations and citations from other authors) received by a publication in the year of its publication or one year after to predict citations two years after publication and later. We calculate the logarithms of the independent variables to manage dispersion and specify two negative binomial regressions and one OLS regression (see Table 5, models 1–3). Model 1 predicts the total number of citations received two years after publication and later. It shows a positive impact of all three types of citations. The influence of unjustified reciprocal citations is the weakest. Model 2 predicts the number of citations received from other authors (i.e., both types of reciprocal citations are removed from the dependent variable) two years after publication and later. The model shows that only citations by other authors have an influence. Justified and unjustified reciprocal citations are not significant. Model 3 is an OLS regression and uses the journal-weighted number of citations received as the dependent variable. The model shows that unjustified and justified reciprocal citations have no influence. Citations from other authors have a positive influence but this effect is very small.

Overall, the results provide limited evidence that unjustified reciprocal citations have a weaker or even a negative effect on future citations compared to other types of citations. Considering the total citation count or citations from other authors suggests a positive impact of citations by others (models 1 and 2). Using the journal-weighted number of citations received indicates also a very small effect of citations from other authors (model 3).

Reciprocal citations of both types are statistically insignificant except for predicting the total number of citations received. In that case, we found a weaker impact of reciprocal citations, especially the unjustified reciprocal citations. Thus, we only have very weak support for H2b.

5. Discussion and conclusion

5.1. Summary

This paper tackled the issue of reputation manipulation in innovation and knowledge exchange communities. Based on the concept of reputation-based reward systems (Dasgupta and David, 1994), reputation has been characterized as enabling selective incentives to compensate the contributors. We distinguished between explicit and implicit representation of reputation. The

[8] The assumptions for Poisson regressions in models 1 and 2 are not fulfilled because of over-dispersion. Therefore, we calculated negative binominal regressions.

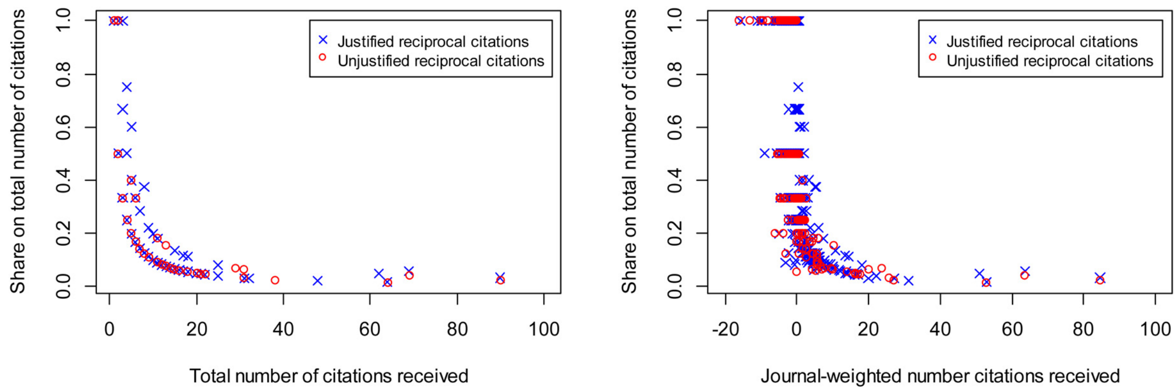

Fig. 2. Share of justified and unjustified reciprocal citations in relation to citations received; left: total number of citations received; right: journal-weighted number of citations received.

explicitness of reputation is a crucial element for online communities. In contrast to implicit reputation, explicitness of reputation reduces participants' effort keep track of others' reputations and allows communities to grow. It increases reliability and minimizes inconsistencies. However, the explicitness of reputation causes difficulties because it opens opportunities for manipulation. While reputation-based reward systems relying on implicit reputation seem to have developed mechanisms to reduce opportunities for manipulation (e.g., social mechanisms counteract self-praise and similar attempts to inflate one's own reputation), explicit reputation seems to be more vulnerable to manipulation attempts.

We investigated science as an example of an explicit reputation-based reward system to get a better understanding of manipulation in such systems. In science, citations represent an explicit measure of contributions' value. They render authors' contributions widely visible and permanent (Merton, 1988; Stephan, 1996). Our analysis of a citation tree extracted from Thomson Reuters' Web of Science that shows strong indicators of manipulation in form of self- and reciprocal citing. Specifically, the data show that authorship, beyond content, predicts citation relationships (H1a and H1b). Publications that receive unjustified self-citations (i.e., citations with overlaps in authors but no overlaps of content) are less often cited than publications that receive justified self-citations (i.e., citations with overlaps in both authors and content) or citations from other authors. This indicates that the publications receiving unjustified self-citations are of inferior quality (H2a). We found only very limited evidence that reciprocal citing is associated with inferior quality (H2b). A reason might be the small number of citations that could be labeled as unjustified reciprocal citations.

5.2. Contribution to theory

This paper contributes to the extant literature in several ways. First, we extend our understanding of communities by introducing the concept of explicit reputation. Communities are inherently affected by problems of collective action (Allen, 1983; Ostrom, 2000). Reputation can mitigate these problems. Information and communication technology plays a crucial role not only for more efficient communication but also for enabling the codification of information on participants' contributions, thus making reputation information explicit. The explicitness of reputation increases its value as a form of information. Knowledge exchange communities, such as answers.com and Stack Overflow, are examples for the power of explicit reputation as an important driver of participation. In addition to these entirely new communities, some offline communities are entering the online world. This step allows them to migrate from implicit to explicit reputation-based reward systems. Examples are the community of haute cuisine chefs, who innovate collectively by exchanging knowledge (Fauchart and von Hippel, 2008; di Stefano et al., 2014). Much of their communication has moved online (Fauchart and von Hippel, 2008). In such cases, the step into the online world also simplifies identification of the originator of a contribution.

Second, this paper makes a fundamental step towards a holistic theory of communities of knowledge exchange and innovation. Although there is ample research on communities (e.g., Franke and Shah, 2003; Lüthje et al., 2005; Jeppesen and Frederiksen, 2006; Hew and Hara, 2007; Faraj et al., 2011; Antorini et al., 2012; Hienerth et al., 2014; Levine and Prietula, 2014), no unified theory has been developed. The private-collective model of innovation in open source software (von Hippel and von Krogh, 2003) could be considered the basis for a unified theory. Private-collective innovation relies on several kinds of selective incentives (Lerner and Tirole, 2002; Wasko and Faraj, 2005; Henkel, 2006; Roberts et al., 2006). However, understanding the motivations of these mechanisms and their functioning requires some elaboration, specifically the role of reputation. We contribute to this by enhancing the understanding of reputation and manipulation in the light of selective incentives (Oliver, 2013; Olson, 1965). We argue that selective incentives are accompanied by motivating the community participants to manipulate their reputations. Social control mechanisms which prohibit the manipulation of own reputation are limited, especially if reputation is explicit.

Third, the paper extends the research stream of the economics of science (Aghion et al., 2009; Stephan, 1996; Stephan, 2012), especially the branch of misconduct (see Broad, 1981; Martin, 2013). It shows empirically that manipulation can take the forms of self- and reciprocal citing. Self-citing is a well-known phenomenon (Glänzel et al., 2004; Glänzel and Thijs, 2004), but it is not deeply discussed in the context of misconduct. Reciprocal citing, in general, has rarely been examined. By addressing these rather soft forms of manipulation, this paper offers new insights to the discussion of the abilities and limits of self-governance in the scientific system (Dasgupta and David, 1994; Martin, 2012). It could be argued that the current scientific system works despite these manipulation activities. However, it is unlikely that this system provides scientists with optimal incentives. Manipulation activities of some participants demotivate the honest contributors because their returns shrink. Consequently, they might reduce their efforts or even adopt manipulation behavior themselves.

5.3. Limitations and future research

This paper has several limitations. The most important one is related to the challenge of distinguishing manipulation from non-manipulation (unjustified from justified self- or reciprocal

Table 4
Prediction of the number of citations received by the types of citations received early (unjustified and justified self-citations and citations from other authors) (self-citing).

	Model 1: negative binominal prediction of citations received (2nd year and later)		Model 2: negative binominal prediction of citations received from other authors (2nd year and later)		Model 3: OLS regression prediction of journal-weighted citations received (2nd year and later)	
	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*t*-value
Unjustified self-citations	0.27*** (0.01)	31.2	−0.07*** (0.01)	−6.5	−0.24*** (0.11)	−2.3
Justified self-citations	0.43*** (0.01)	64.4	−0.01*** (0.01)	−0.8	0.74*** (0.07)	11.0
Citations from other authors	1.01*** (0.00)	237.1	1.33*** (0.00)	280.1	0.58*** (0.01)	48.9
Constant	−0.20*** (0.01)	−27.6	−0.74*** (0.01)	−84.4	−14.90*** (0.08)	−182.0
Log-likelihood	−42,717.6		−36,208.6			
Adjusted R^2					.101	
F					936.3***	
n	24,912		24,912		24,912	

* $p < .05$.
*** $p < .001$.

Table 5
Prediction of the number of received citations by the type of citations received early (unjustified and justified reciprocal citations and citations from other authors) (reciprocal citing).

	Model 1: negative binominal prediction of citations received (2nd year and later)		Model 2: negative binominal prediction of citations received from other authors (2nd year and later)		Model 3: OLS regression prediction of journal-weighted citations received (2nd year and later)	
	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*t*-value
Unjustified reciprocal citations	0.28*** (0.06)	4.8	0.03*** (0.06)	0.5	0.01*** (0.04)	0.0
Justified reciprocal citations	0.34*** (0.04)	8.6	−0.01*** (0.04)	−0.3	0.01*** (0.02)*	−0.0
Citations from other authors	0.91*** (0.01)	133.3	1.00*** (0.01)	157.7	0.01*** (0.00)	4.7
Constant	−0.08*** (0.01)	−7.7	−0.22*** (0.01)	−21.4	−0.55*** (0.00)	−15.9
Log-likelihood	−25,536.0		−24,696.6			
Adjusted R^2					.001	
F					7.513***	
n	14,632		14,632		14,632	

*** $p < .001$.

citations). It is impossible to capture an author's motive precisely; that is, to identify citing for the purposes of manipulation or citing to acknowledge content, or other reasons such as citing publications by authorities to raise the awareness of one's own work. There is a natural information asymmetry because authors cannot credibly reveal their motivations to cite.[9] We controlled for content using keywords, journal, and abstract to minimize this problem. However, we can only use these controls as a heuristic and not as a precise indicator to distinguish manipulation from non-manipulation. At the level of individual citations, they are insufficient to identify single instances of manipulation. However, at the aggregate level, they provide an indicator for distinguishing citations that are more likely associated with a content-based motive from those citations that are probably driven by the intention to manipulate.

Another important limitation is that this study ignores the specific characteristics of scientific communities. Authors in highly specialized or small research communities are more likely to self-cite and to reciprocally cite one another. We argue that our content control minimizes a potential bias in specialized communities. Nevertheless, future research could test the influence of specialization and other characteristics of scientific communities. In addition to size and degree of specialization of a scientific community, it would be interesting to investigate the influence of research funding on manipulation behavior.

Future research should try to develop a deeper understanding of reputation-based reward systems and manipulation. Understanding these systems requires consideration of their dynamics such as self-reinforcement (i.e., positive feedback). Dynamics are crucial to infer a system's stability limits (i.e., how much manipulation the community can bear) and its resilience (i.e., the community's ability to recover from shocks). The current paper does not capture the dynamics of the system, but we would encourage future research in this direction. Speculating about the effects of self-reinforcing dynamics emphasizes the negative impact of manipulation. The dynamics might amplify the impact of self- and reciprocal citation behavior on the reputation-based reward system. As citations attract further citations (Merton, 1968; Merton, 1988), manipulated citations cause proper (i.e., non-manipulated) citations (see Fowler and Aksnes, 2007). Thus, self-reinforcement increases the disturbance by manipulations out of proportion. In this context, our results suggest that the journal-weighted citation count could be a more robust measure. It discriminated strongly between unjustified and justified self-citations and thus can be used for discounting the impact of self-citations. Future research should build on that and examine different weightings and measures to detect manipulation.

[9] Note that this information asymmetry makes self-citing and reciprocal citing tempting because authors can always claim to have honest motives.

Another possibility for future research is to use citation networks to investigate other forms of manipulation in science. This approach makes it possible to address more serious forms of scientific misconduct than self- and reciprocal citing. For instance, the highly criticized practice of self-plagiarism (see Martin, 2013) could be analyzed on a larger scale. The first hint of deliberate self-plagiarism might be a lack of a citation relationship between two publications despite overlapping content *and* overlapping authorship.

Future research should address the differentiation between implicit and explicit reputation in reputation-based reward systems. Developing a better understanding of the advantages of explicit reputation requires direct comparisons with implicit reputation-based reward systems. Although we characterized explicit reputation as desirable, it is unclear whether it actually has a positive effect on performance. It might increase competition and lead to performance enhancement. Alternatively, the competition might demotivate parts of the community and reduce performance.

5.4. Implications for policymakers and community governance

This paper informs policymakers and managers dealing with community-based innovation and knowledge exchange. Reputation-based reward systems determine the contributions and actions that are rewarded or are penalized. The design of the system should be aimed at generating the desired behavior, that is, the optimal level of participants' effort. Therefore, we offer three recommendations for policymakers and managers.

First, we encourage policymakers and managers of implicit reputation-based reward systems to transform reputation into an explicit representation. The conceptual argumentation developed in this paper highlights the advantages of the explicitness of reputation. Maintaining communities with large numbers of members becomes more feasible. Disputes over members' reputations and the status hierarchy are reduced, which leads to more contributions (see Kilduff et al., 2016). However, as noted earlier, switching to an explicit representation of reputation is not always desirable because it could increase the competition among community members. Thus, switching to an explicit representation of reputation should be conducted carefully.

Second, we recommend that designers and managers of innovation and knowledge exchange communities consider the possibilities for manipulation. Identification of manipulation attempts and activities should be a major goal in community governance. Automatic mechanisms to detect manipulation attempts are technologically possible. Especially communities of open source software development (Osterloh and Rota, 2007; von Krogh and von Hippel, 2006; von Krogh et al., 2012) can profit from this recommendation. In these communities, the topic of unethical behavior in form of inappropriate reuse of software code has been observed (Sojer et al., 2014). Platforms hosting open source software projects, such as SourceForge and GitHub, can counteract such behaviors by defining rules that govern reputation and consider manipulation attempts.

Third, policymakers in the area of science are encouraged to isolate self-citations automatically by marking them and making them distinct. The Web of Science offers a rudimentary way to discount self-citations. However, only the self-citations of one author can be identified; the self-citations of co-authors are treated as normal citations. This is inconsistent with characterizing citations as a measure of the contribution to the collective action (Carley et al., 2013). A radical move would be to prohibit self-citing. However, this would likely shift reciprocal citing, which is more difficult to detect. More importantly, forbidding self-citing would facilitate self-plagiarism. Authors who deliberately reuse their own work in a redundant or contradictory way avoid citations between their publications (García-Romero and Estrada-Lorenzo, 2014). The general prohibiting of self-citations would make it easier to disguise such cases. Isolating self-citations by marking seems to be the most satisfactory solution. Authors would be able to refer to their own prior work but the reputation-based reward system could discount for these self-citations.

Acknowledgments

I thank Rudolf Mathey for developing the graph database. I thank Christina Raasch, Henry Sauermann, and Tim Schweisfurth for valuable comments and suggestions on earlier versions of the paper. I thank Saskia Carstensen, Clemens Harten, and Christian Matt for their help with the data collection. The research was supported by funding from the DFG SFB 768 "Managing cycles in innovation processes—Integrated development of product service systems based on technical products," subproject C5.

Appendix A. Robustness tests

We replicate the tests for H2a by weakening the measure for controlling content. Our strategy to test the quality of self-cited publications relies on distinguishing between justified and unjustified self-citations. In the paper, we use only the keywords as the indicator for content and ignored the abstract and the journal. To show that this operationalization is uncritical, we replicate the cal-

Table A1
Replication of prediction of the number of citations received by the type of citations received early (unjustified and justified self-citations and citations from other authors) (self-citing).

	Model 1: negative binominal prediction of citations received (2nd year and later)		Model 2: negative binominal prediction of citations received from other authors (2nd year and later)		Model 3: OLS regression prediction of journal-weighted citations received (2nd year and later)	
	Coefficient	*z*-value	Coefficient	*z*-value	Coefficient	*t*-value
Unjustified self-citations	0.14*** (0.02)	6.1	−0.06*** (0.03)	−2.2	−1.77*** (0.36)	−4.9
Justified self-citations	0.47*** (0.01)	79.1	−0.03*** (0.01)	−4.3	0.51*** (0.05)	9.5
Citations from other authors	1.01*** (0.00)	241.3	1.33*** (0.00)	281.1	0.58*** (0.01)	48.6
Constant	−0.23*** (0.01)	−30.5	−0.74*** (0.01)	−84.1	−14.90*** (0.08)	−181.9
Log-likelihood	−42,528.4		36,208.6			
Adjusted R^2						.101
F						930.7***
n	24,912		24,912		.101	

*** $p < .001$.

culation by using overlap in journal, keywords, or the abstract as indicator for content overlap to operationalize this differentiation. Thus, if one indicator matches, content overlap is assumed. This is very conservative and leads to only a very few unjustified self-citations. We find 653 unjustified self-citations, 22,227 justified self-citations, and 84,280 citations from other authors.

The results from the regressions (Table A1) show a similar pattern as we have shown previously (cf., Table 4). All types of citations have a positive impact on the total number of citations. The impact of unjustified self-citations is the weakest (model 1). For citations from other authors, both types of self-citation have a negative influence (model 2). (The justified self-citations are insignificant in the original analysis, cf., Table 4.) The measure of the journal-weighted number of citations received shows even clearer the negative influence of unjustified self-citations (model 3).

References

Aghion, P., David, P.A., Foray, D., 2009. Science, technology and innovation for economic growth: linking policy research and practice in 'STIG Systems'. Res. Policy 38 (4), 681–693.
Aksnes, D.W., 2003. A macro-study of self-citation. Scientometrics 56 (2), 235–246.
Allen, R.C., 1983. Collective invention. J. Econ. Behav. Organ. 4 (1), 1–24.
Antorini, Y.M., Albert, M., Muniz, J., Askildsen, T., 2012. Collaborating with customer communities: lessons from the Lego Group. Sloan Manage. Rev. 53 (3), 73–79.
Axelrod, R., 1984. The Evolution of Cooperation. Basic Books, New York, NY.
Belenzon, S., 2012. Cumulative innovation and market value: evidence from patent citations. Econ. J. 122 (559), 265–285.
Bolton, G., Greiner, B., Ockenfels, A., 2013. Engineering trust: reciprocity in the production of reputation information. Manage. Sci. 59 (2), 265–285.
Bonzi, S., Snyder, H.W., 1991. Motivations for citation: a comparison of self citation and citation to others. Scientometrics 21 (2), 245–254.
Boudreau, K.J., Lakhani, K.R., 2015. Open disclosure of innovations, incentives and follow-on reuse: theory on processes of cumulative innovation and a field experiment in computational biology. Res. Policy 44 (1), 4–19.
Broad, W.J., 1981. The publishing game: getting more for less. Science 211 (4487), 1137–1139.
Carley, S., Porter, A.L., Youtie, J., 2013. Toward a more precise definition of self-citation. Scientometrics 94 (2), 777–780.
Charness, G., Masclet, D., Villeval, M.C., 2014. The dark side of competition for status. Manage. Sci. 60 (1), 38–55.
Dahlander, L., Magnusson, M.G., 2005. Relationships between open source software companies and communities: observations from Nordic firms. Res. Policy 34 (4), 481–493.
Dasgupta, P., David, P.A., 1994. Toward a new economics of science. Res. Policy 23 (5), 487–521.
Deci, E.L., Ryan, R.M., 1985. Intrinsic Motivation and Self-Determination in Human Behavior. Plenum, New York.
Dellarocas, C., 2006. How often should reputation mechanisms update a trader's reputation profile? Inform. Syst. Res. 17 (3), 271–285.
di Stefano, G., King, A.A., Verona, G., 2014. Kitchen confidential? Norms for the use of transferred knowledge in gourmet cuisine. Strategic Manage. J. 35 (11), 1645–1670.
Enders, W., Hoover, G.A., 2004. Whose line is it? Plagiarism in economics. J. Econ. Liter. 42 (2), 487–493.
Faraj, S., Jarvenpaa, S.L., Majchrzak, A., 2011. Knowledge collaboration in online communities. Organ. Sci. 22 (5), 1224–1239.
Faraj, S., Johnson, S.L., 2011. Network exchange patterns in online communities. Organ. Sci. 22 (6), 1464–1480.
Fauchart, E., von Hippel, E., 2008. Norms-based intellectual property systems: the case of French chefs. Organ. Sci. 19 (2), 187–201.
Fehr, E., Gächter, S., 2000. Fairness and retaliation: the economics of reciprocity. J. Econ. Perspect. 14 (3), 159–181.
Firth, D., 1993. Bias reduction of maximum likelihood estimates. Biometrika 80 (1), 27–38.
Fleming, L., Waguespack, D.M., 2007. Brokerage, boundary spanning, and leadership in open innovation communities. Organ. Sci. 18 (2), 165–180.
Fowler, J.H., Aksnes, D.W., 2007. Does self-citation pay? Scientometrics 72 (3), 427–437.
Franke, N., Shah, S., 2003. How communities support innovative activities: an exploration of assistance and sharing among end-users. Res. Policy 32 (1), 157–178.
Gagné, M., Deci, E.L., 2005. Self-determination theory and work motivation. J. Organ. Behav. 26 (4), 331–362.
García-Romero, A., Estrada-Lorenzo, J.M., 2014. A bibliometric analysis of plagiarism and self-plagiarism through déjà vu. Scientometrics 101 (1), 381–396.
Garfield, E., 1986. Do Nobel Prize winners write citation classics? Essays Inform. Scientist 9 (23), 182–187.
Garfield, E., Welljams-Dorof, A., 1992. Of Nobel class: a citation perspective on high impact research authors. Theoret. Med. 13 (2), 117–135.
Glänzel, W., Debackere, K., Thijs, B., Schubert, A., 2006. A concise review on the role of author self-citations in information science, bibliometrics and science policy. Scientometrics 67 (2), 263–277.
Glänzel, W., Thijs, B., 2004. The influence of author self-citations on bibliometric macro indicators. Scientometrics 59 (3), 281–310.
Glänzel, W., Thijs, B., Schlemmer, B., 2004. A bibliometric approach to the role of author self-citations in scientific communication. Scientometrics 59 (1), 63–77.
Gouldner, A.W., 1960. The norm of reciprocity: a preliminary statement. Am. Sociol. Rev. 25 (2), 161–178.
Hardin, G., 1968. The tragedy of the commons. Science 162 (3859), 1243–1248.
Heinze, G., Ploner, M., 2003. Fixing the nonconvergence bug in logistic regression with SPLUS and SAS. Comput. Methods Programs Biomed. 71 (2), 181–187.
Henkel, J., 2006. Selective revealing in open innovation processes: the case of embedded Linux. Res. Policy 35 (7), 953–969.
Hertel, G., Niedner, S., Herrmann, S., 2003. Motivation of software developers in open source projects: an internet-based survey of contributors to the Linux kernel. Res. Policy 32 (7), 1159–1177.
Hew, K.F., Hara, N., 2007. Knowledge sharing in online environments: a qualitative case study. J. Am. Soc. Inform. Sci. Technol. 58 (14), 2310–2324.
Hienerth, C., von Hippel, E., Jensen, M.B., 2014. User community vs. producer innovation development efficiency: a first empirical study. Res. Policy 43 (1), 190–201.
Hutter, K., Füller, J., Hautz, J., Bilgram, V., Matzler, K., 2015. Machiavellianism or morality: which behavior pays off in online innovation contests? J. Manage. Inform. Syst. 32 (3), 197–228.
Hyland, K., 2003. Self-citation and self-reference: credibility and promotion in academic publication. J. Am. Soc. Inform. Sci. Technol. 54 (3), 251–259.
Jeppesen, L.B., Frederiksen, L., 2006. Why do users contribute to firm-hosted user communities? The case of computer-controlled music instruments. Organ. Sci. 17 (1), 45–63.
Kostoff, R.N., 1998. The use and misuse of citation analysis in research evaluation. Scientometrics 43 (1), 27–43.
Kilduff, G.J., Willer, R., Anderson, C., 2016. Hierarchy and its discontents: status disagreement leads to withdrawal of contribution and lower group performance. Organ. Sci. 27 (2), 373–390.
Lacetera, N., Zirulia, L., 2011. The economics of scientific misconduct. J. Law Econ. Organ. 27 (3), 568–603.
Lange, D., Lee, P.M., Dai, Y., 2011. Organizational reputation: a review. J. Manage. 37 (1), 153–184.
Lanzolla, G., Frankort, H.T.W., 2016. The online shadow of offline signals: which sellers get contacted in online B2B marketplaces? Acad. Manage. J. 59 (1), 207–231.
Lerner, J., Tirole, J., 2002. Some simple economics of open source. J. Ind. Econ. 50 (2), 197–234.
Levine, S.S., Prietula, M., 2014. Open collaboration for innovation: principles and performance. Organ. Sci. 25 (5), 1414–1433.
Lüthje, C., Herstatt, C., von Hippel, E., 2005. User innovators and "local" information: the case of mountain biking. Res. Policy 34 (6), 951–965.
Martin, B.R., 2012. Does peer review work as a self-policing mechanism in preventing misconduct: a case study of a serial plagiarist. In: Mayer, T., Steneck, N. (Eds.), Promoting Research Integrity in a Global Environment. World Scientific Publishing, pp. 97–114.
Martin, B.R., 2013. Whither research integrity? Plagiarism, self-plagiarism and coercive citation in an age of research assessment. Res. Policy 42 (5), 1005–1014.
Merton, R.K., 1968. The Matthew effect in science. Science 159 (3810), 56–63.
Merton, R.K., 1973. The Sociology of Science: Theoretical and Empirical Investigation. University of Chicago Press, Chicago, IL.
Merton, R.K., 1988. The Matthew effect in science. II. Cumulative advantage and the symbolism of intellectual property. Isis 79 (4), 606–623.
Milinski, M., Semmann, D., Bakker, T.C.M., Krambeck, H.-J., 2001. Cooperation through indirect reciprocity: image scoring or standing strategy? Proc. Roy. Soc. Lond. B: Biol. Sci. 268 (1484), 2495–2501.
Milinski, M., Semmann, D., Krambeck, H.-J., 2002. Reputation helps solve the 'tragedy of the commons'. Nature 415 (6870), 424–426.
Moreno, A., Terwiesch, C., 2014. Doing business with strangers: reputation in online service marketplaces. Inform. Syst. Res. 25 (4), 865–886.
Nowak, M.A., Sigmund, K., 2005. Evolution of indirect reciprocity. Nature 437 (7063), 1291–1298.
Oliver, P., 2013. Selective incentives. In: Snow, D.A., Porta, D.D., Klandermans, B., McAdam, D. (Eds.), The Wiley-Blackwell Encyclopedia of Social and Political Movements. Wiley, Malden, MA.
Olson, M., 1965. The Logic of Collective Action: Public Goods and the Theory of Groups. Harvard University Press, Cambridge, MA.
Osterloh, M., Rota, S., 2007. Open source software development: just another case of collective invention? Res. Policy 36 (2), 157–171.
Ostrom, E., 2000. Collective action and the evolution of social norms. J. Econ. Perspect. 14 (3), 137–158.
Ostrom, E., 2007. Challenges and growth: The development of the interdisciplinary field of institutional analysis. J. Inst. Econ. 3 (3), 239–264.
Paisley, W., 1990. An oasis where many trails cross: the improbable cocitation networks of a multidiscipline. J. Am. Soc. Inform. Sci. 41 (6), 459–468.
Phelan, T.J., 1999. A compendium of issues for citation analysis. Scientometrics 45 (1), 117–136.

Piezunka, H., Dahlander, L., 2015. Distant search, narrow attention: how crowding alters organizations' filtering of suggestions in crowdsourcing. Acad. Manage. J. 58 (3), 856–880.
Posner, R.A., 2000. An economic analysis of the use of citations in the law. Am. Law Econ. Rev. 2 (2), 381–406.
Powell, W.W., 1990. Neither market nor hierarchy: network forms of organization. In: Staw, B.M., Cummings, L.L. (Eds.), Research in Organizational Behavior. JAI Press Inc, Greenwich, CT, pp. 295–336.
Roberts, J.A., Hann, I.-H., Slaughter, S.A., 2006. Understanding the motivations, participation, and performance of open source software developers: a longitudinal study of the Apache projects. Manage. Sci. 52 (7), 984–999.
Rullani, F., Haefliger, S., 2013. The periphery on stage: the intra-organizational dynamics in online communities of creation. Res. Policy 42 (4), 941–953.
Schemmann, B., Herrmann, A.M., Chappin, M.M., Heimeriks, G.J., 2016. Crowdsourcing ideas: involving ordinary users in the ideation phase of new product development. Res. Policy 45 (6), 1145–1154.
Schubert, A., Glänzel, W., Thijs, B., 2006. The weight of author self-citations: a fractional approach to self-citation counting. Scientometrics 67 (3), 503–514.
Snyder, H., Bonzi, S., 1998. Patterns of self-citation across disciplines 1980–1989. J. Inform. Sci. 24 (6), 431–435.
Sojer, M., Alexy, O., Kleinknecht, S., Henkel, J., 2014. Understanding the drivers of unethical programming behavior: the inappropriate reuse of internet-accessible code. J. Manage. Inform. Syst. 31 (3), 287–325.
Stam, W., 2009. When does community participation enhance the performance of open source software companies? Res. Policy 38 (8), 1288–1299.
Stephan, P.E., 2012. How Economics Shapes Science. Harvard University Press, Cambridge, MA.
Stephan, P.E., 1996. The economics of science. J. Econ. Liter. 34 (3), 1199–1235.
Stephan, P.E., 2010. The economics of science. In: Arrow, K.J., Intriligator, M.D. (Eds.), Handbook of the Economics of Innovation, vol. I. North-Holland, Amsterdam, pp. 217–273.
Trivers, R.L., 1971. The evolution of reciprocal altruism. Quart. Rev. Biol. 46 (1), 35–57.
van Raan, A.F., 2008. Self-citation as an impact-reinforcing mechanism in the science system. J. Am. Soc. Inform. Sci. Technol. 59 (10), 1631–1643.
von Hippel, E.A., von Krogh, G., 2003. Open source software and the "private-collective" innovation model: issues for organization science. Organ. Sci. 14 (2), 209–223.
von Krogh, G., Haefliger, S., Spaeth, S., Wallin, M.W., 2012. Carrots and rainbows: motivation and social practice in open source software development. MIS Quart. 36 (2), 649–676.
von Krogh, G., Spaeth, S., Lakhani, K.R., 2003. Community, joining, and specialization in open source software innovation: a case study. Res. Policy 32 (7), 1217–1241.
von Krogh, G., von Hippel, E., 2006. The promise of research on open source software. Manage. Sci. 52 (7), 975–983.
Wasko, M.M., Faraj, S., 2005. Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. MIS Quart. 29 (1), 35–57.
West, J., Lakhani, K.R., 2008. Getting clear about communities in open innovation. Ind. Innov. 15 (2), 223–231.
White, H.D., Wellman, B., Nazer, N., 2004. Does citation reflect social structure? Longitudinal evidence from the "Globenet" interdisciplinary research group. J. Am. Soc. Inform. Sci. Technol. 55 (2), 111–126.
Yoganarasimhan, H., 2013. The value of reputation in an online freelance marketplace. Market. Sci. 32 (6), 860–891.