

Original citation:

Aaltonen, Aleksi and Lanzara, Giovan Francesco. (2015) Building governance capability in online social production : insights from Wikipedia. *Organization Studies*, 36 (12). pp. 1649-1673.

Permanent WRAP url:

<http://wrap.warwick.ac.uk/75774>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions.

This article is made available under the Creative Commons Attribution 3.0 (CC BY 3.0) license and may be reused according to the conditions of the license. For more details see: <http://creativecommons.org/licenses/by/3.0/>

A note on versions:

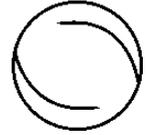
The version presented in WRAP is the published version, or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: publications@warwick.ac.uk

warwick**publications**wrap

highlight your research

<http://wrap.warwick.ac.uk>



Building Governance Capability in Online Social Production: Insights from Wikipedia

Organization Studies
2015, Vol. 36(12) 1649–1673

© The Author(s) 2015

Reprints and permissions: 
sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/0170840615584459

www.egosnet.org/os



Aleksi Aaltonen

Warwick Business School, The University of Warwick, UK

Giovan Francesco Lanzara

University of Bologna, Italy

Abstract

This article investigates a form of governance that makes online social production possible. Drawing on the concepts of capability and routine, we develop a dynamic, process-oriented view that departs from past research focused on static comparative analysis. We theorize that online social production systems develop a collective governance capability to steer the process of integrating distributed knowledge resources to the production of value. Governance mechanisms emerge from individual and collective learning that is made possible by new technology, and they evolve over time, as routines are developed to respond to new problems faced by a growing production system. Using Wikipedia as a paradigmatic example of online social production, we characterize governance as an evolving, enabling and embedded process and discuss implications for a dynamic theory of governance.

Keywords

capability, governance, knowledge-based view, social production, Wikipedia

Introduction

Today, thousands of weakly connected individuals are able to collaborate and deliver consistent online output without the price or corporate system governing the activity. Benkler (2006) describes the phenomenon as social production, i.e. an alternative form of organizing production that is facilitated by contemporary digital technology. In this article, we analyse the development of Wikipedia to understand the form of governance that makes social production possible in the online environment.

Corresponding author:

Aleksi Aaltonen, Warwick Business School, The University of Warwick, Coventry, CV4 7AL, UK.

Email: aake@iki.fi

We define *online social production* as an internet-mediated activity that is directed toward specific outputs in the absence of individual contracts and managerial authority to control and coordinate human action. Linux, Wikipedia, OpenStreetMap and many other projects have demonstrated that such efforts can successfully compete with corporate production in many different industries. Online social production creates innovative products through new ways of organizing and harnessing knowledge resources at remarkably low coordination costs. Yet, framing online social production systems in terms of organizational economics provides only a limited view to the phenomenon. The systems have shown that new governance arrangements can convert widely dispersed knowledge into a resource that is largely inaccessible through market or corporate arrangements (Benkler, 2006; Shirky, 2008).

Whatever advantages online social production may have, large-scale production systems cannot avoid governance problems that pertain to how to steer the integration of dispersed knowledge resources and how to coordinate such activities to the purpose of creating common value. Millions of minuscule contributions do not fall together into complex products such as an operating system or an encyclopedia by accident. They must be assessed, selected, adapted and oftentimes rejected, that is, work must be governed to the purpose of creating value. We claim that a more thorough understanding of how governance is carried out in these systems is critical to unleashing their collective potential for value creation in business and human affairs.

The way we approach governance departs from perspectives based organizational economics (Williamson, 1996). Rather than focusing on control and on the transaction cost efficiency of different institutional arrangements (Benkler, 2006, pp. 59–63; Demil & Lecocq, 2006; Lee & Cole, 2003; O'Mahony & Ferraro, 2007), we contend that governance is an evolving phenomenon in online social production. The configuration of mechanisms that makes joint action possible has to support different functions over time as governance adapts to the growing complexity and relative maturity of a system and its product. In contrast, current accounts tend to privilege a static comparative analysis of distinct governance modes by assessing their transaction cost efficiency, which fails to capture the dynamic, evolutionary character of governance that stands out in many longitudinal analyses of online social production systems (e.g. Aaltonen & Kallinikos, 2013; Cornford, Shaikh, & Ciborra, 2010; Fitzgerald, 2006; Forte, Larco, & Bruckman, 2009; Halfaker, Geiger, Morgan, & Riedl, 2012).

To develop a more dynamic view of governance, we adopt a capability-based perspective (Nelson & Winter, 1982; Teece, Pisano, & Shuen, 1997; Zollo & Winter, 2002; Winter, 2003). We theorize that governance in online social production rests on the progressive development of a collective capability to integrate highly distributed knowledge resources and direct them to the joint production of value. A *capability* is what an individual, organization or other type of collective arrangement can actually do, which is expressed as learned patterns of repetitive behaviour, that is, *routines* (Jacobides & Winter, 2012; Nelson & Winter 1982; Pentland & Feldman, 2005; Winter, 2003). We assume that governance itself can be regarded as a type of capability; it is the capability to progressively design and implement mechanisms to control and coordinate joint production (Argyres, Felin, Foss, & Zenger, 2012). We define *collective governance capability* as the capability of a collective arrangement to steer a production process and an associated interaction system. It is collective in the sense that the governance capability does not depend on any particular individual or a group of individuals in the system. Importantly, the collective governance capability does not have to be supported by a nexus of contracts (that is, a firm) and may not be present at the outset, when the system is *in statu nascenti*. We ask therefore the following research question:

How does a collective governance capability to create and maintain value emerge and evolve in online social production?

To answer the question, we study routines that have made Wikipedia governable at different times. We look into how specific routines or sets of interacting routines may perform as governance mechanisms in such a large, distributed system. Therefore, by *governance mechanism* we refer either to a single routine or to the interaction of sets of routines by which a specific governance aim is attained. The online encyclopedia is one of the most successful examples of online social production to date and revolves on processes that are typical of such systems, the most basic one being the process of integrating knowledge from dispersed sources to produce encyclopedia articles and to solve various governance problems that appear along the way (Grandori, 1997, 2001; Michailova & Foss, 2009). Adopting a longitudinal, process-oriented case study approach (Huber & Van de Ven, 1995; Pettigrew, 1990; Yin, 2014), we show that governance is based on mechanisms that are: *evolving* over time rather than fixed, becoming more articulated as they integrate new knowledge about the system's own operations; *enabling*, that is, primarily facilitating learning and knowledge integration and sharing, and largely *embedded* in the interaction and production system rather than residing in a separate administrative body. These attributes cannot be understood separate from the digital technology that absorbs much of the organizing that is accomplished in online social production. The findings contribute both to an improved understanding of how governance happens in online social production systems and to governance theory by offering a more dynamic, process-oriented view.

Theoretical Framework

According to von Hayek (1945), governance issues arise when knowledge relevant to the production of value exists fragmented among different actors. Although von Hayek mainly focused on the market as a mechanism to govern knowledge exchanges, the view equally applies to online social production. To put it simply, it is far from clear how highly distributed knowledge can be integrated and steered to a coherent collective output in the online environment (Foss, 2007; Michailova & Foss, 2009; Grandori, 2013; Langlois & Foss, 1999). Online social production systems are characterized by various conditions of imperfect knowledge: *where* knowledge useful for the production task is located and *how* it is distributed; *how* to coordinate the work and skills of thousands of dispersed agents interacting over the internet; and *how* to maintain a complex production system over time, protecting and enhancing the value it creates.

The problem of governance in online social production

The specific features of online social production (Aaltonen & Kallinikos, 2013; Benkler, 2006; Niederer & van Dick, 2010) make governance a distinct problem with respect to market-based coordination of dispersed knowledge (von Hayek, 1945) or to coordination by a central authority like the firm or the state (Coase, 1937; Kogut & Zander, 1992; North, 1990). There are no employment contracts or formal role structure to fall back on, and resources are relatively free to flow in and out of the system, so that the governing agents may not even know whom they are actually governing. Basic conditions for firm-based governance mechanisms and structures do not simply exist (Foss, Husted, & Michailova, 2010; Grandori, 2001; Kogut & Zander, 1996). More generally, hierarchy is not effective in governing highly distributed knowledge (Grandori, 2009; von Hayek, 1945); market contracts are no good for recruiting and coordinating resources due to high friction on exchanges (Demil & Lecocq, 2006), and communitarian bonds tend to be thin and volatile in these systems, because identities are mostly unknown or ill-defined, especially in the early phases of system evolution (Faraj, Jarvenpaa, & Majchrzak, 2011).

Many online social production systems have nevertheless developed effective governance mechanisms that maintain access to resources and facilitate the integration of knowledge and its conversion to high-quality output. This has often required developing multiple heterogeneous mechanisms that enable responding to problems that arise within an evolving system. Steering a handful of people around a nascent idea poses for sure different problems than holding together thousands of individuals working on a mature product. The rapid growth and fluid character of online social production systems (Faraj et al., 2011; Schreyögg & Sydow, 2010) make the change of governance mechanisms particularly salient, inviting to look for a dynamic perspective.

Current governance theory tends to privilege the comparative assessment of different institutional arrangements that support social coordination and administrative regulation, but underplays the dynamic processes underlying governance (Bevir, 2013; Kooiman, 2003; Ostrom, 1990; Rhodes, 2007; Stoker, 1998; Williamson, 1996). This has led to studies emphasizing either social bonds and shared norms that support a communitarian form of (self)governance in online social production (Lee & Cole, 2003; von Krogh & von Hippel, 2006; Watson et al., 2005) or a bureaucratic apparatus that emerges to cope with increasing administrative complexity (Crowston & Howison, 2005; Joyce, Pike, & Butler, 2013). Other accounts posit hybrid combinations of market, community and hierarchy that try to balance the composite, heterogeneous requirements of online social production (Demil & Lecocq, 2006; O'Mahony & Ferraro, 2007; Shah, 2006). All these accounts tend to fall back on a static view of governance anchored to discrete institutional forms, failing to capture the process by which new governance mechanisms are developed in online social production.

Collective governance capability

The idea that organizations embody capabilities that cannot be reduced to their individual members has its origin and the primary domain of application in the knowledge-based view (KBV) of the firm (Dosi, Faillo, & Marengo, 2008; Grant, 1996a; 1996b; Helfat et al., 2007; Kogut & Zander, 1992, 1996; Langlois & Foss, 1999). Firms integrate knowledge into organizational capabilities to accomplish complex tasks and create value by transforming inputs into outputs (Grant, 1996b; Jacobides & Winter, 2012; Nelson & Winter, 1982). Capabilities, in other words, embody the organizational know-how 'that enables individuals and organizations to perform [...] activities' (Dosi et al., 2008, p. 1165).

The KBV has not theorized governance with comparable conceptual detail to organizational economics, but the idea that governance is a distinct capability that is learned and developed by integrating knowledge, and is therefore an evolving asset, has been acknowledged in recent debates on the theoretical convergence of capabilities and organizational economics (Argyres et al., 2012; Jacobides & Winter, 2012). Once a capability has been developed and consolidated as an asset, it may become itself an object of governance: therefore a governance capability may ultimately include the governance of other capabilities (Argyres et al., 2012; Winter, 2003). Knowledge integrated into a governance capability can be understood as knowledge about controlling and coordinating a large-scale distributed system, which is closely related to but different from production-oriented capabilities such as writing encyclopedia articles or developing software. Capabilities are usually found anchored to traditional organizational forms, but it is reasonable to assume that they emerge in other types of collective arrangements as well. An online social production system can be defined as a bundle of capabilities for doing things (Dosi et al., 2008) – just like an organization – although the circumstances for capability-building are different.

In a firm context, studies based on KBV can assume the existence of certain resources and managerial capacity to steer them toward corporate objectives (Grant, 1996a). The starting point

for online social production is different: there is no managerial capacity, potential resources exist widely distributed in society, and their pattern of distribution is unknown at the outset. Possibilities for configuring the system *ex ante* are therefore much more limited and, whatever plans founders may make, those cannot be managerially enforced. A successful online social production system must learn to execute new and progressively more complex tasks and to steer production toward shifting objectives as a result of experience from ongoing multiple activities. This is reflected in the consolidation of new routines.

Routines and mechanisms

Capabilities cannot be observed directly but are embodied in bundles of interrelated routines that encode practical experience and knowledge (Pentland & Feldman, 2005; Nelson & Winter, 1982; Winter, 2003). Routines are the building blocks of capabilities and express patterned and repetitive behaviour that is learned. Yet, they are not only procedures to get things done, but, as Dosi and colleagues (2008, p. 1167) point out, routines also have a regulative dimension. They are expressions of rule-driven behaviour that is essential to governance or, in a more generic sense, to the ordering and predictability of activities (March & Simon, 1958). The governance side of routines is often deeply embedded in the system of activities and directly enacted in the execution, to an extent that it can be difficult to distinguish governance from the immediate production or problem-solving side of routine operations.

When a routine or an interacting set of routines comes to play a specific regulative function in a system, it is then said to perform as a governance mechanism. The term *mechanism* evokes here the functional or rule-based dimension of governance rather than its agency component, and it is widely used in organizational economics and policy sciences (Bevir, 2013; Foss, Husted, & Michailova, 2010; Grandori, 1997; Kooiman, 2003; Rhodes, 2007; Williamson, 1996). To better appreciate the extent to which governance is embedded in the system of routines and operations, we prefer to make the following distinction: we use 'routine' whenever we focus on agency and behavioural patterns that build up capabilities, and use 'mechanism' whenever we want to emphasize the functions routines play in the governance system.

A new capability can be said to be fully learned only when it is enacted in practice. Observable behavioural patterns can be studied empirically and, as our analysis shows, can also become an object of governance interventions. Therefore, the learning of new capabilities in online social production is revealed by new routines that are collectively enacted. If the collective arrangement learns specific routines to jointly accomplish a task, and the enacted routines can be empirically observed, then a collective capability for doing things is revealed.

In synthesis, our conceptual framework builds on capability theory to account for the nature of governance in large-scale online social production systems. We theorize governance as a collective capability to steer the process of collecting and integrating distributed knowledge resources to the creation and maintenance of valuable products and services. Such a capability rests upon multiple and heterogeneous sets of routines that respond to specific governance problems arising with the growing size and complexity of the system. The framework is intended to help us capture the dynamic, process-like character of governance rather than its static, structural properties.

Methodology, Data and Descriptive Statistics

Wikipedia originated in January 2001 as an experimental spinoff from another encyclopedia project based on a more traditional editorial approach. The English Wikipedia currently consists of over 4.5 million articles as a result of hundreds of millions of voluntary contributions over the

years.¹ Every month, tens of thousands of contributors flock in and out of the system that remains highly organized and able to steer its own development despite the considerable fluidity of its resources.

Wikipedia is produced using an open-source technological platform called MediaWiki. The platform makes it possible to use a web browser to write web pages, link them together and add other features using a simple markup code. The pages are divided into 29 mutually exclusive domains known as namespaces that designate the type of content presented on a page.² The main namespace contains encyclopedia articles, while others are dedicated to discussions (talk pages), policies and guidelines, user profiles, and so forth. A link between two pages can represent a semantic association or it can construct a functional relationship between items. For instance, linking an article page to a page describing a content category assigns the article to that category; an editorial operation can be justified by leaving a talk page comment with a link to an applicable policy. The pages are stored in a version control system that keeps track of all changes made to Wikipedia.

In order to describe and analyse the development of such a massive system, we combine two sets of research operations into a sequential multi-methods study (Edmondson & McManus, 2007; Venkatesh, Brown, & Bala, 2013). First, we calculate and briefly report descriptive statistics on how the overall shape of Wikipedia has changed over time. The statistics point to different periods in the evolution of the system and its governance, but they cannot tell how or why Wikipedia has changed. Second, using the picture provided by the statistics as a general frame of reference, we proceed to the analysis of governance routines as they appear in the system. More specifically, we provide a longitudinal interpretation of evolving governance mechanisms based on the assessment of historical versions of Wikipedia pages. The qualitative analysis provides an explanation of the observed quantitative pattern including meta-inferences combining different types of evidence (Venkatesh et al., 2013).

Descriptive statistics

Wikimedia Foundation has made the raw data of different Wikipedia sites freely available for research purposes. We use the `enwiki-20100130-pages-meta-history.xml.bz2` file that contains the full text of nearly 350 million edits on 19 million different pages from the inception of the English Wikipedia on 16 January 2001 to 30 January 2010. The content of the edits range from fixing a small typographic error to adding (or removing) an entire page. As such, the edits represent the enactment of basic routines embedded in the technological platform. We process the data into 108 monthly snapshots from January 2001 to December 2009, summarizing the data by individual contributors. This reveals a known fact about the differentiation of Wikipedia contributors. An active member of Wikipedia community may make several small edits on many different pages and be active across the namespaces, while a topical expert may concentrate his contributions to fewer but more sizeable edits on just a few encyclopedia articles (Kittur, Chi, Pendleton, Suh, & Mytkowicz, 2007). To get a longitudinal view of contributor differentiation, we perform a cluster analysis to observe the appearance of contributor groupings based on editing behaviour (see also Aaltonen & Lanzara, 2011). The technique identifies natural groupings in multidimensional data by minimizing intra-group differences and by maximizing inter-group differences between the cases (see Appendix 1 for the description of statistical procedures).

Figure 1 plots the number of monthly contributors and observed contributor clusters in the English Wikipedia over the nine-year observation period. The number of contributors across all namespaces varies from six in January 2001 to 191,138 in March 2007. Consistent with Halfaker and colleagues (2012, p. 665), we identify a slow initiation period followed by a period of extremely

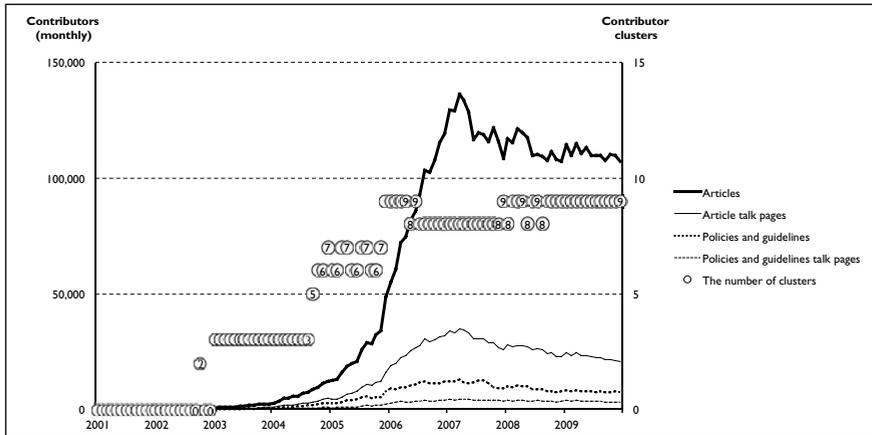


Figure 1. The number of monthly contributors and the number of contributor clusters in the English Wikipedia from January 2001 to December 2009.

rapid growth, and, finally, levelling out and a slight decline. In the first phase, there is only a minimal differentiation of contributors into clusters. The second phase of exponential growth is characterized by increasing differentiation of contributors, while the number of clusters stabilizes in the third phase. The statistics provide only a very rough depiction of a complex system, but they certainly suggest that, whatever governance mechanisms have been in place, they have had to deal with dramatically different circumstances over the years. The statistics do not suggest that there are clear-cut boundaries between the periods and, in this sense, our periodization is best understood as an analytical scaffold that helps to focus qualitative analysis on issues that have been critical at different times.

Qualitative analysis

The qualitative analysis is carried out as a theoretical narrative, that is, a theory-driven description of a developmental process. Polkinghorne (1988, p. 13) describes narrative as ‘the fundamental scheme for linking individual human actions and events into aspects of understandable composite’, to which we add the idea of writing the plot around theoretical constructs rather than individual actors. The narrative is made of a sequence of interconnected events and states that fit together as a coherent story from the perspective of capability-building. The approach draws from the idea that longitudinal organizational research (Kimberly, 1976; Miller & Friesen, 1982; Thomson, Plumridge, & Holland, 2003) can be fruitfully combined with theory building to generate a retrospective interpretation of a process (Hernes, 2008).

Our narrative is constructed as a highly selective representation of empirical material that advances the plot of capability-building. The analysis aims to establish a conceptual structure in a concrete storyline and, at the same time, to present empirical findings in a way that can be inspected, tested and eventually contested. For this purpose, the narrative form helps to trace shifting relationships between events and patterns in an unfolding process much more flexibly than a variable-based framework would allow (Abell, 2004; Pentland, 1999, p. 711). The blending of online observations and descriptive statistics together with a pinch of ‘disciplined imagination’ (Weick, 1989) makes it possible to account for routines that signal capabilities addressing

governance problems at different phases of Wikipedia evolution. For these reasons, we find a theoretical narrative to be an apt way to analyse how subsequent patterns and configurations emerge from previous ones through the interplay of conflicting forces and requirements (Schreyögg & Sydow, 2010).

Theoretical Narrative: Building Governance Capability in Wikipedia

In this section, we combine the general frame of reference provided by the descriptive statistics with online observations to construct a theoretical narrative of how capabilities addressing specific governance problems emerge in Wikipedia. Wikipedia pages used as illustrative examples are indexed as [1, 2, 3, ...] and listed in Appendix 2 including a full online address. The narrative outlines the building of governance mechanisms supporting a specific kind of value. The value of Wikipedia, as defined in the early vision of the founders [1], is to be a free, open and easily accessible knowledge source for the largest possible number of users; also, it is to be a reliable, peer-approved resource and the collective outcome of a large crowd of independent editors.³

The early years: attracting and integrating distributed knowledge resources

Figure 1 shows that the website initially attracted only a handful of contributors, and consequently its content grew slowly for the first few years. There were less than 3000 monthly active user accounts in the end of 2003, and the lack of content was an overarching problem for the encyclopedia-to-be. Any educated person had at this stage topical knowledge that could add value to an empty encyclopedia, which explains the rationale for the all-out, indiscriminate resource mobilization in the beginning. These circumstances are reflected vividly in an early stance toward rules in Wikipedia, for one of the oldest principles of Wikipedia is known as ‘Ignore all rules’. Taken literally, such a paradoxical rule would seem to undermine any attempt to develop governance mechanisms. The principle reads in its original form:

If rules make you nervous and depressed, and not desirous of participating in the Wiki, then ignore them and go about your business. (“Wikipedia:Ignore all rules” page [2] on 17 April 2002)

A more accurate interpretation of the ‘Ignore all rules’ principle is probably to take it as an expression that rule following should not hinder people’s motivation to contribute content. This, rather than denying governance as such, calibrates it to present circumstances and to the paramount concern for resource mobilization. The principle is perhaps also an admission that early contributors often faced situations in which any extant rule would not make sense. People were therefore lured to do whatever they saw fit, while nobody knew exactly how to manage the development of a peer-produced encyclopedia. The standard organizational solution of centralizing decision-making was not applicable for a number of reasons. Wikipedia had no managerial capacity and very limited resources to start with: an open-source technological platform, few volunteers and no content. Relevant knowledge on how to control and coordinate activities did not exist because no similar product had been created online, which resembles the problem of governance identified by von Hayek (1945). The venture harnessed an open-source technological platform and internet connectivity to bootstrap the encyclopedia. The initial assets made it possible to enact a few basic routines to work on individual encyclopedia articles: a writing routine, version control routine, reverting routine and a discussion routine. These routines have been at the core of content production and the basis for capability-building in Wikipedia from the beginning.

Anybody can enact the *writing routine* by clicking the 'Edit' tab on top of a Wikipedia page, which activates a writing tool. The person can then edit the page content and click 'Save page' to store a new version of the article that is published immediately. This is the most rudimentary production routine and all that is required to make an edit to the online encyclopedia. The writing routine is supported by a version control technology that imports coordination mechanisms used in software development practices. *Version control routines* avoid splitting an article edited by multiple contributors into parallel versions, provide a complete trace of edits from different contributors, and allow reverting a page to any of its earlier versions. *Reverting routine* is one of the most important editorial routines in Wikipedia and, importantly, is equally available to all contributors. It is used to undo mistakes, to reject inappropriate or substandard edits, and also to try forcing one's own viewpoint and vandalizing the system. However, the openly available edit history makes it possible for individual contributors to hold each other accountable and to revert inappropriate reverts when needed.

Incoming contributors brought in topical knowledge but also different ideas about the encyclopedia itself and how to collaborate with others, which were then refined in the process of jointly producing articles. In addition to the writing, version control and reverting routines, direct interaction between contributors is facilitated by a dedicated talk page attached to each article. Talk pages allow enacting the same writing, version control and reverting routines as article pages but differ in their content. They are used to post short personal communications about the associated article and, therefore, make it possible to enact a *discussion routine* focusing on a specific piece of content. The writing, version control, reverting and discussion routines provide key means by which Wikipedia sourced its first contributions from people who became, for one reason or another, attracted to the experiment.

The early contributors brought in a mix of individual capabilities that became available to online social production, but no collective capability at the level of encyclopedia production can be inferred at this stage. The article and its talk page provided a minimal framework for peer production, allowing people to build on one another's contributions. In the process, contributors became gradually more skilful at integrating diverse knowledge into encyclopedia articles by repeatedly enacting the basic routines and by resolving problems on talk pages with their peers. Contributors working on contentious topics had to also learn how to build epistemic consensus through the resolution of controversies and the merging of multiple viewpoints into a coherent article (see Table 1).

As long as the number of contributors remained sufficiently small, the version control, reverting and discussion routines provided a control and coordination mechanism that was sufficient for most situations. Governance was embedded in interaction between individuals, revolved on individual encyclopedia articles, and was supported by a few generic principles such as 'Ignore all rules'. Each article could be developed in relative isolation from others, allowing the system to accommodate a variety of individual knowledge, skills and interests without surfacing too many conflicts and contradictions. The circumstances were also highly conducive to the development of new knowledge, which mainly occurred in the context of individual articles and among small groups of contributors. Minimal and unstable as this learning process may have initially been, it gave rise to a common editing practice and set off a common intent.

The growth of complexity: the emergence of a collective governance capability

During the early years, a central problem was to attract individual knowledge and skills at reporting on different topics. Resources were found diffused in society as 'cognitive surplus' (Shirky, 2010), i.e. an unexploited potential that could be converted into productive use on the new technological platform. Figure 1 shows that Wikipedia attracted gradually more and more contributors, and in late 2003 the growth started to accelerate rapidly. Accelerating growth signalled the viability

Table 1. Building governance capability in the early days of Wikipedia.

Governance problem	How to attract and integrate distributed knowledge resources?
Examples of routines	<ul style="list-style-type: none"> • Writing routine • Version control routines • Reverting routine • Discussion routine
Capabilities	<p>Capabilities are focused to the production of encyclopedia articles:</p> <ul style="list-style-type: none"> • Individual skills and knowledge in writing on topic • Technological ordering of edits from multiple contributors • Collaborative assessment of edit quality • Discussion focused on article content and its development
Learning	Contributors learn from each other in talk page discussions and by observing reactions to edits
Social structure of capabilities	Capabilities are anchored to small and fluid groupings of contributors and to the technological platform

of the project that was also starting to break into public awareness. Wikipedia aroused more and more people's curiosity, attracting larger and larger numbers of contributors, which put considerable pressure on the early governance mechanisms.

A simple accumulation of content does not yet make a useful encyclopedia product. The epistemic value of an encyclopedia resides in being a coherent body of knowledge organized by a system of categories and cross-references, reliable updates, etc. The question of encyclopedic quality culminated in 2005, when *Nature* compared Wikipedia to *Encyclopædia Britannica* in a widely reported study (Giles, 2005). The study found selected Wikipedia articles surprisingly reliable, which ever since has been used as a testament to the viability of the novel approach to building an encyclopedia. The event was a major public relations victory for Wikipedia, but it also cemented *Encyclopædia Britannica* as the gold standard by which the output of the online social production system would be evaluated. New requirements were consequently found to be important to encyclopedic knowledge: reliability, coverage, consistency, epistemic consensus and frequent updates. The inadequacy of original governance mechanisms became clear as a result of a hoax biography article on a well-known journalist, John Seigenthaler. The article was written in May 2005, but it went unnoticed and uncorrected until September, causing embarrassment to the system that escalated all the way to the mainstream media. New routines emerged as a response to problems.

We use Three-revert rule (3RR) and Featured Article Review (FAR) procedure to illustrate the kinds of developments that describe Wikipedia during this period. 3RR appeared as a result of collective deliberation by a vote in November 2004. The rule established new routines that tackle behaviour instantiated by the writing and reverting routines in a radically open and increasingly popular production system. The common right to immediately edit and revert others' contributions had both enabled the rapid integration of knowledge from a massive audience but also led to unproductive and malicious behaviour. Problems started to appear as the contributor base grew and became more diverse in its views and motivations that could no longer be reconciled by peer discussion. 3RR was particularly designed to curb a phenomenon known as edit warring, that is, a situation in which the joint production of an article turns into antagonistic back-and-forth reversals between different views.

The Three-revert rule (or 3RR) is an official policy which applies to all Wikipedians. The policy states that an editor must not perform more than three reversions, in whole or in part, on a single Wikipedia article within a 24 hour period. (“Wikipedia:Three-revert rule” page [3] on 2 March 2006)

An article subject to an edit war does not converge to a high-quality encyclopedia entry but instead results in a waste of resources and unreliable content. 3RR limits the right to repeatedly enact the reverting routine on the same page and authorizes the use of a blocking routine to enforce the limit. It also refers to further dispute resolution mechanisms in the case a disagreement cannot be solved between the contributors working on the disputed article. The rule is therefore a response to editing behaviour that arose from the particular social and technical configuration of Wikipedia, and from the inability to control such behaviour using earlier governance mechanisms based on peer-interaction. 3RR inscribes the boundaries of the acceptable use of the basic reverting routine.

The Featured Article Review arose from increasing concerns about the lack of editorial standards. A page listing particularly well-developed articles had already appeared as early as November 2001, which developed in the latter half of 2003 into a procedure for nominating, assessing and promoting featured articles. The procedure became eventually established as a sort of production line for flagship articles signalling quality and attention to standards in Wikipedia.

The featured articles are what Wikipedian editors believe to be the best articles in Wikipedia. Prior to being listed here, articles are reviewed at Wikipedia:Featured article candidates for accuracy, neutrality, completeness, and style according to our featured article criteria. (“Wikipedia:Featured articles” page [4] on 1 August 2006)

Any contributor can propose a Featured Article candidate, which triggers a review process during which the article is developed and assessed against common criteria that apply to the entire encyclopedia. Contributors can participate in the editing of the candidate article as usual and, finally, the Featured Article Director decides if there is a strong enough consensus to promote the status of the article. The nomination, assessment and promotion routines employed encode new knowledge on how to carry out quality assurance in a system based on open participation of a diverse population. They constitute a mechanism that aligns the skills and knowledge of an individual contributor with objectives and structures that describe the system in general. The routines that substantiate 3RR and FAR are expressions of a collective governance capability: they were developed inside Wikipedia and are specific to it; their rationale transcends individual articles and is bound to the collective arrangement as a whole, and, finally, they do not produce any content but rather adjust how the contributors should enact the basic production routines. Both 3RR and FAR channel and restrict contributor behaviour, and hence function as governance mechanisms with a broader scope beyond individual articles. More generally, 3RR and FAR can be said to contribute toward a minimal managerial agency that, in the latter case, is also attached to a new role in the system.

We observe how the basic writing, version control and reverting routines that underpin the production of encyclopedia articles assume additional functions to counter new governance problems. They are used to create and integrate knowledge about online social production itself. Contributors discuss content production problems on talk pages and externalize their learning by writing rudimentary policies that are stored as pages in a separate domain (namespace) in the system. Such writing about encyclopedia production itself gave rise to a metatext through which contributors can construct and legitimize new kinds of routines and agency. The metatext gradually took on a more important role in governance vis-a-vis peer interaction as contributors

Table 2. Building collective governance capability in the period of rapid growth.

Governance problem	How to control and coordinate a distributed and rapidly growing production system?
Examples of routines	<ul style="list-style-type: none"> • Three-revert rule (3RR) routines • Featured Article Review (FAR) routines
Capabilities	<p>New capabilities are anchored to the online social production system rather than to individual contributors or small groupings</p> <p>Examples:</p> <ul style="list-style-type: none"> • Capability to control behaviour instantiated by the writing and reverting routines in a radically open system • Capability to improve the quality of articles against a common criteria
Learning	Contributors develop new routines by discussing problems on talk pages and writing metatext; they also learn through the enactment of the new routines
Social structure of capabilities	The enactment of production routines remain widely distributed, but some editorial and administrative agency become more centralized and attached to emerging roles

incorporated their learning into policy pages and then used the metatext to justify actions. The capacity of the same technological platform to simultaneously support production and governance functions curbs structural complexity, since contributors can shift between different types of tasks using the same tool. It also allowed governance to grow from within production and only gradually differentiate as a separate domain of activity. As a result, anybody who is able to participate in the production of encyclopedia content can, at least in principle, also contribute to its governance. Reflecting upon developments during this period, ‘Wikipedia:Editorial oversight and control’ page from early 2007 describes:

The Wikipedia community is largely self-organising, so that anyone may build a reputation as a competent editor and become involved in any role they may choose, subject to peer approval. Individuals often will choose to become involved in specialised tasks, such as reviewing articles at others request, watching current edits for vandalism, or watching newly created articles for quality control purposes, or similar roles. (“Wikipedia:Editorial oversight and control” page [5] on 8 January 2007)

The excerpt confirms important developments in the system. Regular contributors had begun to perceive the need for governance across individual articles and, along with the appearance of routines that take the overall encyclopedia content as their frame of reference, administrative roles such as the Featured Article Director started to emerge. Figure 1 shows differentiation into contributors clusters characterized by different activities (see also Aaltonen & Lanzara, 2011). This resulted both from restrictions on the default rights to operate and from contributors choosing to focus on specific activities that appear in the system. The execution of governance routines remained, however, embedded in or very close to the production of articles. Some new routines belong exclusively to emerging roles, but in general contributors carry out editorial and administrative routines while still being engaged in content production (see Table 2).

The changes during this period mean that the object to be governed is no longer an individual article but the body of encyclopedic knowledge, which reflects the development of a collective capability at controlling and coordinating the production of a complex knowledge artefact. The

capability is manifested in the configuration of content production and governance routines that make the orderly production of value possible.

The age of maturity: maintaining and enhancing the common value

On 4 August 2006, Jimmy Wales argued in his Wikimania conference keynote speech that the main focus of work should move from the production of new content to improving quality [6]. Figure 1 shows that the growth of contributors levelled in early 2007. The work done by a massive contributor base takes place now in a completely different setting from the early days. Content production is governed through a network of routines and, while contributions are still welcome from anyone, there is progressively less content to add and more to maintain in the encyclopedia.

The consolidation of the collective capability and the maturation of encyclopedia content created new problems. The levelling growth and gradual decline in participation seen in Figure 1 indicate that, despite the explosive growth and enthusiasm in building the free encyclopedia, its maintenance is not guaranteed. Accordingly, the governance problem shifts from developing the rules of content production to their efficient enforcement, from routines steering how the basic production routines are enacted to routines that are aimed at strengthening the collective capability and protecting the valuable assets that have accumulated over the years. In the following, we examine new routines that appear to shape how already established governance routines are enacted.

The first example is a policy that governs the use of bots to automate routines. A Wikipedia bot (shorthand for 'robot') is a software-controlled user account that performs repetitive operations with little or no human intervention. They are used for content production and editorial tasks such as making small grammatical corrections and, historically, for importing bulk content from other sources. The use of bots was discouraged for several years. However, when the focus shifted to protecting Wikipedia's assets and collective capability, perceptions about bots and the policy changed. Sophisticated bots offer an efficient way to monitor contributor behaviour and enforce rules. For instance, a user account ClueBot-NG belongs to a bot that uses machine learning techniques to monitor edits for potential vandalism and enacts the reverting routine on vandalized pages [7].

Prior to use, bots must be approved at [Wikipedia:Bots/Requests for approval](#). State there precisely what the bot will do, observe and participate in the discussion, and await authorization from someone in the approvals group. Approval may take up to one week. After approval you may run the bot for a short period so it can be monitored. After the trial the bot will be evaluated by someone in the approvals group and you may ask that the user be marked as a bot by a bureaucrat. ("Wikipedia: Bot policy" page [8] on 17 January 2007)

The bot policy page from early 2007 lists a number of routines intended to govern the deployment of bots: an approval routine, a trial and evaluation routine, and a routine for requesting a user account to be marked as a bot. It is important to note that these routines are not concerned with the actual tasks the bot may perform but govern the deployment of the bot mechanism itself. Therefore, we are witnessing the emergence of routines that are aimed at the system of governance itself, to strengthen the capability to maintain Wikipedia, whereas they are only indirectly related to the production of encyclopedia content.

The second example is the introduction of a flagged revisions feature and a policy to introduce a mandatory review on selected pages before new revisions are published. Wikipedia was founded on the principle of unrestricted participation, so that anybody could write articles, have a say on anything, and get his or her views published. In 2007, however, a flagged revisions feature appears

that allows enforcing new routines to protect the encyclopedia from potentially low-quality contributions. By activating the feature on a page, it is possible to introduce an editorial oversight routine to content production.

The control known as flagged revisions is being rolled out as of 2007. It aims to differentiate the version shown to most readers, from the draft ‘cutting edge’ version being edited, and in the first instance to only show the latter when it has been checked for reasonableness. Flagged revisions are expected to provide a powerful way to prevent most vandalism or poor quality edits from being seen by readers, once it is fully operational. (“Wikipedia: Editorial oversight and control” page [9] on 19 October 2007)

Whenever the feature is applied to an article, it separates the writing routine from publishing a new version of the article. The right to contribute to the common good is therefore separated from the right to be published, and new routines are needed to settle how and by whose decision flagged revisions are published. The flagged revisions routine does not directly concern the substantive quality of encyclopedia content or even how contributors should behave, but it creates a space for new routines to adjust to the effects of the basic production routines. This represents a significant shift with respect to the original principles of the system and its tagline ‘the free encyclopedia that anyone can edit’. Against this background, it is not surprising that it took several years before the community could agree on a policy governing the use of the flagged revisions feature.

Contrary to some media reports, flagged revisions were not enabled in 2009 on the English Wikipedia. A considerably less restrictive version was trialled in 2010–2011. After lengthy discussions, the Wikipedia community agreed a policy to use the feature from 1 December 2012.

(“Wikipedia: Flagged revisions” page [10] on 14 September 2013)

The development of Three-revert rule, Featured Article Review, bot policy, flagged revisions and other routines expresses the rise of a collective capability to govern a complex production system. Many of the important routines were conceived inside Wikipedia and are clearly specific to it. The knowledge embedded in them was largely developed by enacting the same writing, version control, reverting and discussion routines that are used to integrate distributed knowledge into encyclopedia articles. Contributors play an important role in this process, but the online social production system is more than a collection of individuals. Its routines inscribe learning that is collective and cannot be reduced to individual actors. These developments have had implications both for the kind of contributors Wikipedia needs and opportunities for further learning. Today, it is necessary that a contributor becomes familiar with numerous routines before he or she can contribute effectively. Rather than people burning to share their knowledge and experiment with online social production, Wikipedia needs disciplined contributors who are willing to become socialized into a sophisticated editorial process and to learn its rules (see Table 3).

The object being governed is no longer the encyclopedia product but the online social production system itself. This does not mean that earlier routines and governance mechanisms would disappear. Some may, but many of them remain a part of an increasingly dense network of principles (Ignore all rules), rules (3RR), procedures (FAR), policies (bot policy), software instruments and collective resolutions (flagged revisions) that bear upon one another and the basic production routines. The principle ‘Ignore all rules’ still exists today, but it has become very difficult to find a situation in which no existing rule would apply.

Table 3. Building collective governance capability in the age of maturity.

Governance problem	How to protect and maintain the online social production system?
Examples of routines	<ul style="list-style-type: none"> • Bot deployment routines • Flagged revisions routines
Capabilities	<p>New capabilities target the collective governance capability itself</p> <p>Examples:</p> <ul style="list-style-type: none"> • Capability to stabilize capabilities by automating routines • Capability to balance participation and quality in the production of articles
Learning	Contributors are socialized to a regime of principles, rules, procedures, policies, etc.; learning increasingly happens through norms and rules
Social structure of capabilities	The enactment of production routines remain mostly distributed despite some selective restrictions, while a concentrated and structured system of administrative capabilities is established

Governance as an Evolving, Enabling and Embedded Process

The narrative provides a number of insights on how the creation of value is made possible by the emergence of a collective governance capability in online social production. Wikipedia stands for a situation characterized by imperfect knowledge of how to govern a large-scale online production process in the absence of traditional foundations of organizing. The findings convey a picture of governance that departs from current views and can be extended to a variety of online social production systems in which highly distributed agents jointly create knowledge-intensive products. More specifically, we can characterize governance in online social production as an *evolving, enabling* and *embedded* process.

Evolving

The theoretical narrative brings to the fore the dynamic, evolving aspects of governance. The Wikipedia production system undergoes dramatic changes over the nine-year observation period, becoming a very different thing than it was at the beginning. The different size, speed of growth and maturity of the product pose a multiplicity of governance problems to Wikipedia that could initially resort to hardly more than peer-to-peer discussion as a problem-solving mechanism. The pattern resembles, for instance, Linux (Cornford et al., 2010; Lee & Cole, 2003), as both systems began as loose experiments with practically nil resources but have nevertheless succeeded in building a collective governance capability, massive assets and a mature product.

The narrative shows how new problems lead to the formation of new routines and capabilities, which lead to new forms of agency and, consequently, to reconfigurations and different dynamics of the system. This, again, causes further problems to emerge. The tendency towards more organization-like structures and processes, revealed by increasing contributor cluster differentiation in Figure 1 and the identification of new routines acting on earlier ones, is an endogenous learning phenomenon: the system learns to respond to shifting governance problems by developing an array of governance mechanisms that enable and support collective agency. In the process, as new layers of routines are progressively applied to the growing body of existing routines, behaviours are modified and the evolution of capabilities redirected (Jacobides & Winter, 2012; Nickerson & Zenger, 2004). This evolutionary pattern, which Wikipedia epitomizes, characterizes large-scale

online social production systems as they face shifting governance problems throughout their development from inception to maturity.

The theoretical narrative illustrates a feedback that runs between the system's evolving structure, the governance problems faced, the routines developed, the capabilities learned, the agency enabled and, in the end, the further evolution of the system. Finally, by taking its own governance mechanism as the object of new routines, the system can be said to develop a sort of reflexivity. Wikipedia is no longer just a system to produce an encyclopedia but an institution looking to reproduce itself.

New governance mechanisms, however, do not necessarily displace earlier ones. When online social production systems expand, the initial capability to attract and integrate distributed knowledge resources remains important. The original aim must be accomplished within a shifting governance framework that balances openness and quality standards, that is, between feeding variety into the system and at the same time controlling and coordinating behaviour in increasing detail. This is not an easy task (Halfaker et al., 2012; Ortega, 2009). The right balance evolves with the production system, and it must necessarily do so, because the nature and intensity of the contradictory pressures upon the system shift over time. Structure and selectivity at an early stage would have turned potential contributors away and choked the collective experiment, while too much openness and lack of structure at the mature stage would make the system unmanageable, thus putting the value that has been accumulated at risk.

Enabling

In each developmental phase we observe governance mechanisms enabling new types of action, so as to make online social production manageable under ever more complex circumstances. In the nascent phase, when the key problem is to attract distributed resources, a few basic principles and routines embedded into the technological platform make joint production possible. Together with the internet connectivity, the basic routines and the minimal governance empower people to contribute to content production without further training, qualifications or explicit guidance.

By making each contribution readily accessible to anyone, the technological platform generates information that enables all contributors to design further actions and swiftly align their own knowledge and agency to the knowledge and agency of others. This apparently inconspicuous feature turns out to be a powerful governance mechanism, because it channels individual contributions while at the same time allowing for distributed monitoring and instant accountability in the production process. The generalized and diffused observability of the production process contributes significantly to its governability. Governance largely happens *by sight*, that is, by granting the possibility of seeing one another and the common project to all, instead of establishing a centralized and dominant point of view, as in the panopticon.

As the size and complexity of the online social production system grow, the governing mechanisms enabled by the technological platform must be complemented by additional mechanisms that enable knowledge integration at the system level. The new conditions demand that individual actions be partially realigned to meet the system-wide governance challenges. For example, 3RR and FAR enable Wikipedia contributors to produce high-quality articles and speed up the editorial process, so that the growing system does not waste resources in unproductive edit warring or in solving the same problems every time anew. Such realignment, though, is not easy to obtain because, when an online social production system reaches a critical mass of contributors and a large size, tensions emerge between individual agency and the cohesion and efficiency of the system. When the encyclopedia matures, the governance framework must enable the protection and maintenance of the common assets. Accordingly, we notice the expansion of automated routines

that are in charge of many production tasks and enable the large-scale monitoring of the editorial process. In these mature conditions contributors are still enabled to coordinate for collective production, but through governance mechanisms that guide human agency in a more structured fashion than in the early phase.

Embedded

In Wikipedia, governance is embedded to a considerable degree in the interaction and production routines and in the technological platform that absorb much of the organizing that goes on in the production system. Governance is not necessarily encapsulated in separate administrative departments, exclusionary functional roles, or specialized forms of agency that describe formal organizational settings. While online social production systems may selectively deploy such structures, their distinctive feature is the substantial amount of overlap and close intertwining between managerial and productive practices. This is due in part to the technological underpinnings and in part to the organizing principles of commons-based production (Aaltonen & Kallinikos, 2013; Benkler, 2006). Software, therefore, is not simply an instrument of production or an administrative tool, but allows combining the two in new ways. For instance, accountability and monitoring can be directly enacted through the use of software rather than through explicit rules and procedures or human supervision. A similar technological embeddedness can be observed in open source software projects, where organizing and coordination are closely intertwined with the inner workings of the software gear.

Due to the embedded character of governance, all contributors who have access to the software tools made available by the platform and are skilled enough to use them can also become governing agents, albeit later on with differential access to a range of governance actions. However, as the Wikipedia case shows, contributors to production, even when they take on administrative tasks or are granted a special administrative role, do not cease to contribute content (Aaltonen & Lanzara, 2011). Governance is mediated and supported by the same technological platform that mediates and supports productive activities. This, we believe, is a crucial organizing feature inasmuch as it facilitates learning, thus making the building and consolidation of a collective capability at governing the system possible.

Theoretical Implications

The three attributes (evolving, enabling, embedded) expand current conceptions of governance in the literature and have important theoretical implications. Extant theories have mostly focused on the choice of efficient structures to handle (trans)actions and enforce social and economic coordination, but they have seldom dealt with issues concerning dynamic adaptation, learning and capability-building. These, though, are precisely the kind of issues that arise when one tries to understand how online social production systems emerge and are able to outperform traditional arrangements. The main differences between current perspectives and the perspective developed in this article are summarized in Table 4.

To begin with, the possibility that governance could itself be a dynamic, evolving process has largely escaped the current literature. From the perspective of organizational economics, governance choices are based on the comparative assessment of alternative institutional arrangements to mitigate conditions of imperfect knowledge, to regulate behaviour and to minimize transaction costs (Grandori, 2013; Williamson, 1996, 2000). Under this view, governance structures and mechanisms can change only by discrete re-contracting or negotiated redesign, that is, by a shift to another arrangement. Similarly, political governance is usually conceptualized as a mode of

Table 4. Theorizing governance: current perspectives vs. capability-based perspective.

Categories	Current perspectives	Capability-based perspective
Focus	Structures	Processes
Objectives	Cost optimization	Value creation
Mechanisms	Control, regulation, choice	Learning, search, knowledge integration
Logic	Static: system efficiency	Dynamic: system evolvability, generativity
Methods	Static comparative analysis	Longitudinal, process-oriented analysis
Unit of observation	Transactions	Routines
Technology	Information-processing tool	Mediating platform
Architectures	Coherent institutional structures	Combinations of complementary mechanisms

administrative regulation (Bevir, 2013; Ostrom, 1990; Rhodes, 2007) in which the comparative assessment consists in gauging policy effectiveness, regulatory performance and, for instance, distributive fairness. Political agreements and rules can also be changed, but only through bargaining at specific points in time, again implying an exogenous force. These conceptualizations cannot account for the multiple and changing governance mechanisms driving the evolution of a system-in-the-making. The rapidly morphing nature of online social production systems calls for continuous redesign and adaptation of governance mechanisms in order to cope with the shifting conditions of imperfect knowledge. By building on the theoretical programme outlined by Jacobides and Winter (2012) to study the evolution of capabilities, we have used the idea of capability to conceptualize governance as a process driven by feedback linking specific problems to the evolution of capabilities and the dynamic nature of the system.

Second, online social production shows that governance does not just economize on transaction costs but can be at the core of collective pursuits and value creation (Langlois & Foss, 1999; Madhok, 1996). Governance mechanisms can support learning, knowledge integration and sharing from a multiplicity of distributed sources, and they should be assessed with respect to their effectiveness in unearthing and converting new knowledge resources to productive purposes. This is a significant departure from current perspectives on governance, and our study demonstrates how such an assessment can be done based on capability theory. In particular, the case of Wikipedia shows how new kinds of governance mechanisms make it possible to create significant value by harnessing knowledge resources that would be very difficult to seize through a market or corporate system.

Third, the lack of a contractually sanctioned governance framework constitutes initially an advantage in online social production. Systems can put all their resources into production and learning instead of safeguarding the integrity of formal structures (Baldwin, 2007; Benkler, 2006). Online social production does away with a burden that comes with having to maintain contractual relationships as people learn to deal directly with each other through peer-based interaction and informal agreements. Learning comes into the picture both as a *process* by which knowledge is encoded in routines expressing collective governance capabilities and as an *outcome* made possible by the evolving governance framework, for intensive peer-to-peer discussion is certainly a more fertile ground for learning than a managerial command structure. Indeed, the success of online social production stems at least partly from the opportunities for learning it offers to a large number of individuals and from the technological capacity to capture and encode this learning effectively.

Fourth, taking the cue from a number of scholars that have advocated the need for micro-foundations of governance and capability theories (Abell, Felin, & Foss, 2008; Foss, Husted, & Michailova, 2010; Grandori, 2001), we have focused on the ‘cogs and wheels’ of governance rather than on generic structures. A major implication is that the integration of knowledge into governance capabilities need not assume an overarching ‘architectural’ form, as postulated by the KBV of the firm (Grant, 1996b; Jacobides, 2006). Collective and distributed governance can be achieved with minimal mediation by formal structure. More specifically, governance capabilities need not be anchored to stable organizational membership and role structures, or partitioned into specialized administrative bodies, but can be directly embedded in minimal interactive and morphing arrangements. These rely upon an array of routines linking highly variable contributions from distributed agents to the production system. Accordingly, we share the view that theoretical and empirical attention should shift from ideal-type structures to changing combinations of governance mechanisms that complement one another (Grandori, 2013).

Finally, our study highlights the central role of contemporary information technology in supporting new forms of governance. These forms are not reducible to technology, yet technology creates the conditions in which new kinds of governance capabilities can emerge. Indeed, a peculiar feature of online social production systems is that both production and governance are mediated by the same technological platform. The capability to govern complex systems of distributed activities is thus largely embedded in a diffused web of software tools and digital artefacts supported by the internet. Embedding routines in software stabilizes the collective governance capability and frees human resources to solve new problems, but also removes discretion at the level of individual pages and transfers matters further to the domain of policy-making. Recent studies (Kallinikos, 2011; Lanzara & Morner, 2005; Lessig, 2006) have shown that in online social production software tends to assume organizing and governance functions in a functionally simplified form so that they can be handled by the computer. As a consequence, governance of interactions and activities around a complex collective task directly exploits the coordinative properties of the technological platform, thus reducing the need for management in a traditional sense.

Final Remarks

In this article we have theorized how a collective governance capability to create value emerges and evolves in the absence of typical contractual, managerial and communitarian grounds that characterize most organizational settings. Rather than existing from the outset, such a capability appears as a result of endogenous learning in online social production: the system learns to respond to new circumstances by developing new routines that are applied to the growing body of existing routines. Linux, Wikipedia, OpenStreetMap and similar online projects can be said to have developed a collective governance capability precisely because they can carry out a complex set of routines that support the creation and maintenance of value. Governance can be characterized as an evolving, enabling and embedded process driven by individual and collective learning that is made possible by contemporary information technology. Technological platforms have a twofold function in the process: they provide both the production tools and the means by which the system and individual agents can learn about production and its governance.

However, due to its explorative, single-case character, our study has a number of limitations that may constitute further research opportunities. First, by choosing to focus on routines as determinants of governance capabilities, we have overlooked related aspects such as roles and other structural elements, and, more generally, the regime of rights that supports online social production systems. For sure, research on the formal and legal aspects of online governance would nicely complement the present study. Second, as Foss (2007) and Grandori (2013) have argued, the

characteristics of governance (simple/complex; hierarchical/communitarian; procedural/interactive; centralized/distributed) may depend on the nature of knowledge that needs to be integrated (codified/non-codified; formal/informal; imitable/unique). Different online social production systems may therefore show significant variations in the types of governance capabilities, due to the distinctive features of the knowledge involved. However, we claim that the overall process of governance capability-building shows a similar phenomenology for the reasons illustrated above. Finally, to overcome the limitations of a single case study more process-based analyses should be conducted on different online social production systems so that the validity of the analysis advanced here could be comparatively tested and variations and singularities assessed.

Despite its limitations, our work brings to attention governance issues that arise in online social production when a variety of dispersed knowledge contributions must be directed and assembled into a valuable and coherent output without curtailing diversity, shunning participation and hampering social learning. Particularly, it shows how in digital environments governance is ‘re-fashioned’, assuming dynamic forms that combine elements belonging to different media and look remarkably different from the governance structures theorized in current accounts.

Acknowledgements

This article results from research that originated within the former TIGAIR (The Information Growth and Internet Research Group) at the London School of Economics and Political Science. We wish to express our gratitude to Anna Grandori, Jannis Kallinikos, Carsten Sørensen, Philipp Genschel, Attila Marton, José-Carlos Mariategui, Maha Shaikh, Aapo Kyrölä and Erkkä Ryyänen and to many other colleagues who have discussed earlier versions of the article in a number of workshops, research seminars and conferences. We also thank the London School of Economics and Political Science for providing the computational resources needed to produce the descriptive statistics. Finally, we thank the Senior Editor and the four anonymous reviewers for patiently helping us across the troubled waters of difficult revisions.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Notes

1. “Wikipedia:Size of Wikipedia” page on 7 July 2014, see http://en.wikipedia.org/w/index.php?title=Wikipedia:Size_of_Wikipedia&oldid=615924147
2. “Wikipedia:Namespace” page on 30 November 2014, see <http://en.wikipedia.org/w/index.php?title=Wikipedia:Namespace&oldid=636063293>
3. In the literature we find different ways of assessing value in Wikipedia. See, for instance, Priedhorsky, Chen and Lam (2007), who measure value by the number of words read or viewed by readers.

References

- Aaltonen, A., & Kallinikos, J. (2013). Coordination and learning in Wikipedia: Revisiting the dynamics of exploitation and exploration. *Research in the Sociology of Organizations*, 37, 161–192.
- Aaltonen, A., & Lanzara, G.F. (2011). *Governing complex social production in the internet: the emergence of a collective capability in Wikipedia*. Paper presented at the Symposium ‘A Decade in Internet Time’, University of Oxford, 21–24 September.
- Abell, P. (2004). Narrative explanation: An alternative to variable-centered explanation? *Annual Review of Sociology*, 30, 287–310.
- Abell, P., Felin, T., & Foss, N. J. (2008). Building micro-foundations for the routines, capabilities, and performance links. *Managerial and Decision Economics*, 29, 489–502.

- Argyres, N., Felin, T., Foss, N., & Zenger, T. (2012). Organizational economics of capability and heterogeneity. *Organization Science*, 23, 1213–1226.
- Baldwin, C. Y. (2007). Where do transactions come from? Modularity, transactions, and the boundaries of firms. *Industrial and Corporate Change*, 17(1), 155–195.
- Benkler, Y. (2006). *The wealth of networks*. New Haven, NJ: Yale University Press.
- Bevir, M. (2013). *A theory of governance*. Berkeley: University of California Press.
- Coase, R. H. (1937). The nature of the firm. *Economica* (New Series), 4(16), 368–405.
- Cornford, T., Shaikh, M., & Ciborra, C. (2010). Hierarchy, laboratory and collective: Unveiling Linux as innovation, machination and constitution. *Journal of the Association of Information Systems*, 11, 809–837.
- Crowston, K., & Howison, J. (2005). Hierarchy and centralization in free and open source software team communications. *Knowledge, Technology and Policy*, 18(4), 65–85.
- Demil, B., & Lecocq, X. (2006). Neither market nor hierarchy nor network: The emergence of bazaar governance. *Organization Studies*, 27, 1447–1466.
- Dosi, G., Faillo, M., & Marengo, L. (2008). Organizational capabilities, patterns of knowledge accumulation and governance structures in business firms: An introduction. *Organization Studies*, 29, 1165–1185.
- Edmondson, A. C., & McManus, S. E. (2007). Methodological fit in management field research. *Academy of Management Review*, 32, 1155–1179.
- Faraj, S., Jarvenpaa, S. L., & Majchrzak, A. (2011). Knowledge collaboration in online communities. *Organization Science*, 22, 1224–1239.
- Fitzgerald, B. (2006). The transformation of open source software. *MIS Quarterly*, 30, 587–598.
- Forte, A., Larco, V., & Bruckman, A. (2009). Decentralization in Wikipedia governance. *Journal of Management Information Systems*, 26, 49–72.
- Foss, N. J. (2007). The emerging knowledge governance approach. *Organization*, 14, 39–52.
- Foss, N. J., Husted, K., & Michailova, S. (2010). Governing knowledge sharing in organizations: Levels of analysis, governance mechanisms, and research directions. *Journal of Management Studies*, 47, 455–482.
- Giles, J. (2005). Internet encyclopaedias go head to head. *Nature*, 438, 900–901.
- Grandori, A. (1997). Governance structures, coordination mechanisms and cognitive models. *Journal of Management and Governance*, 1, 29–42.
- Grandori, A. (2001). Neither hierarchy nor identity: Knowledge governance mechanisms and the theory of the firm. *Journal of Management and Governance*, 5, 381–399.
- Grandori, A. (2009). Polyarchic governance and the growth of knowledge. In N. J. Foss & S. Michailova (Eds.), *Knowledge governance: Process and perspectives* (pp. 81–107). Oxford: Oxford University Press.
- Grandori, A. (2013). *Epistemic economics and organization: Forms of rationality and governance for a wiser economy*. London: Routledge.
- Grant, R. M. (1996a). Prospering in dynamically-competitive environments: Organizational capability as knowledge integration. *Organization Science*, 7, 375–387.
- Grant, R. M. (1996b). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(Winter Special Issue), 109–122.
- Halfaker, A., Geiger, R. S., Morgan, J. T., & Riedl, J. (2012). The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *American Behavioral Scientist*, 57, 664–688.
- Helfat, C., Finkelstein, S., Mitchell, W., Peteraf, M., Sing, H., Teece, D., & Winter, S. (2007). *Dynamic capabilities: Understanding strategic change in organizations*. Malden, MA: Blackwell Publishing.
- Hernes, T. (2008). *Understanding organization as process: Theory for a tangled world*. New York: Routledge.
- Huber, G. P., & Van de Ven, A. (Eds.) (1995). *Longitudinal field research methods: Studying processes of organizational change*. Thousand Oaks, CA: SAGE Publications.
- Jacobides, M. G. (2006). The architecture and design of organizational capabilities. *Industrial and Corporate Change*, 15(1), 151–171.
- Jacobides, M. G., & Winter, S. G. (2012). Capabilities: Structure, agency, and evolution. *Organization Science*, 23, 1365–1381.

- Joyce, E., Pike, J. C., & Butler, B. S. (2013). Rules and roles vs. consensus: Self-governed deliberative mass collaboration bureaucracies. *American Behavioral Scientist*, 57, 576–594.
- Kallinikos, J. (2011). *Governance through technology*. Basingstoke, UK: Palgrave Macmillan.
- Kimberly, J. R. (1976). Issues in the design of longitudinal organizational research. *Sociological Methods and Research*, 4, 321–348.
- Kittur, A., Chi, E., Pendleton, B. A., Suh, B., & Mytkowicz, T. (2007). Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. Paper presented at the Alt.CHI 2007, San Jose, CA.
- Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3, 383–397.
- Kogut, B., & Zander, U. (1996). What firms do? Coordination, identity and learning. *Organization Science*, 7, 502–518.
- Kooiman, J. (2003). *Governing as governance*. London: SAGE Publications.
- Langlois, R. N., & Foss, N. J. (1999). Capabilities and governance: The rebirth of production in the theory of economic organization. *Kyklos*, 52, 201–218.
- Lanzara, G. F., & Morner, M. (2005). Artifacts rule! How organizing happens in open source software projects. In B. Czarniawska & T. Hernes (Eds.), *Actor-network theory and organizing*. Copenhagen, Denmark: Liber and Copenhagen Business School Press.
- Lee, G. K., & Cole, R. E. (2003). From a firm-based to a community-based model of knowledge creation: The case of the Linux kernel development. *Organization Science*, 14, 633–649.
- Lessig, L. (2006). *Code: And other laws of cyberspace, version 2.0*. New York: Basic Books.
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Cybernetics and Control Theory*, 10, 707–710.
- Madhok, A. (1996). The organization of economic activity: Transaction costs, firm capabilities, and the nature of governance. *Organization Science*, 7, 577–590.
- March, J. G., & Simon, H. A. (1958). *Organizations*. New York: John Wiley & Sons.
- Michailova, S., & Foss, N. J. (2009). Knowledge governance: Themes and questions. In N. J. Foss & S. Michailova (Eds.), *Knowledge governance: Perspectives, processes and problems*. Oxford, UK: Oxford University Press.
- Miller, D., & Friesen, P. H. (1982). The longitudinal analysis of organizations: A methodological perspective. *Management Science*, 28, 1013–1034.
- Navarro, G. (2001). A guided tour to approximate string matching. *ACM Computing Surveys*, 33(1), 31–88.
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Cambridge, MA: Belknap Press of Harvard University Press.
- Nickerson, J. A., & Zenger, T. R. (2004). A knowledge-based theory of the firm: The problem-solving perspective. *Organization Science*, 15, 617–632.
- Niederer, S., & van Dick, J. (2010). Wisdom of the crowd or technicity of content? Wikipedia as a sociotechnical system. *New Media and Society*, 12, 1368–1387.
- North, D. (1990). *Institutions, institutional change and economic performance*. Cambridge, UK: Cambridge University Press.
- O'Mahony, S., & Ferraro, F. (2007). The emergence of governance in an open source community. *Academy of Management Journal*, 50, 1079–1106.
- Ortega, F. (2009). *Wikipedia: A quantitative analysis*. PhD thesis. Universidad Rey Juan Carlos, Madrid, Spain.
- Ostrom, E. (1990). *Governing the commons*. Cambridge, UK: Cambridge University Press.
- Pentland, B. T. (1999). Building process theory with narrative: From description to explanation. *Academy of Management Review*, 24, 711–724.
- Pentland, B. T., & Feldman, M. S. (2005). Organizational routines as units of analysis. *Institutional and Corporate Change*, 14, 793–815.
- Pettigrew, A.M. (1990). Longitudinal field research on change: theory and practice. *Organization Science*, 1, 267–292.
- Polkinghorne, D. A. (1988). *Narrative knowing and the human sciences*. Albany: Albany State of New York Press.

- Priedhorsky, R., Chen, J., & Lam, S. K. (2007). Creating, destroying, and restoring value in Wikipedia. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work* (pp. 259–268). New York, NY: ACM.
- Rhodes, R. A. W. (2007). Understanding governance: Ten years on. *Organization Studies*, 28, 1243–1264.
- Schreyögg, G., & Sydow, J. (2010). Organizing for fluidity? Dilemmas of new organizational forms. *Organization Science*, 21, 1251–1262.
- Shah, S. K. (2006). Motivation, governance, and the viability of hybrid forms of open source software development. *Management Science*, 52, 1000–1014.
- Shirky, C. (2008). *Here comes everybody: The power of organizing without organizations*. New York: Penguin.
- Shirky, C. (2010). *Cognitive surplus: Creativity and generosity in a connected age*. London: Penguin Books.
- Stoker, G. (1998). Governance as theory: Five propositions. *International Social Science Journal*, 50(155), 17–28.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18, 509–533.
- Thomson, R., Plumridge, L., & Holland, J. (2003). Editorial. *International Journal of Social Research Methodology*, 6, 185–187.
- Venkatesh, V., Brown, S. A., & Bala, H. (2013). Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems. *MIS Quarterly*, 37, 21–54.
- von Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, 35, 519–530.
- von Krogh, G., & von Hippel, E. (2006). The high promise of research on open source software. *Management Science*, 52, 975–983.
- Watson, R. T., Boudreau, M.-C., Greiner, M., Wynn, D., York, P., & Gul, R. (2005). Governance and global communities. *Journal of International Management*, 11, 125–142.
- Weick, K. E. (1989). Theory construction as disciplined imagination. *Academy of Management Review*, 14, 516–531.
- Williamson, O. E. (1996). *The mechanisms of governance*. Oxford, UK: Oxford University Press.
- Williamson, O. E. (2000). The new institutional economics: Taking stock, looking ahead. *Journal of Economic Literature*, 38, 595–613.
- Winter, S. G. (2003). Understanding dynamic capabilities. *Strategic Management Journal*, 24, 991–995.
- Yin, R. K. (2014). *Case study research: Design and methods* (5th edition). Thousand Oaks, CA: SAGE Publications.
- Zollo, M., & Winter, S. (2002). Deliberate learning and the evolution of dynamic capabilities. *Organization Science*, 13, 339–351.

Author biographies

Aleksi Aaltonen is an Assistant Professor at Warwick Business School. He holds a PhD in Information Systems from the London School of Economics and Political Science. He is broadly interested in technologies of organizing, and has written about business metrics, digital artefacts, data-driven practices in mobile advertising, and governance in social production. His publications have appeared in *MIS Quarterly*, *Management Science*, *Journal of Information Technology* and *Research in the Sociology of Organizations*.

Giovan Francesco Lanzara is Professor of Organization Studies at the University of Bologna, Department of Political and Social Sciences. He has been a Visiting Professor at the London School of Economics and Political Science from 2006 to 2012, in the Department of Management, Innovation and Information Systems Group. In his research work and publications he has made contributions at the intersection of three different fields – organization studies, information systems, and political science – with a specific focus on practice studies and innovation in organizational and institutional settings.

Appendix I

The construction of descriptive statistics

The raw database file `enwiki-20100130-pages-meta-history.xml.bz2` contains the full text of nearly 350 million edits on 19 million different pages from the inception of the English Wikipedia on 16 January 2001 to 30 January 2010. We processed the file in three steps to produce descriptive statistics. A *contributor* is defined as a Wikipedia user who has made at least one edit to encyclopedia articles or any other namespace.

First, we transformed the original XML-formatted data into a tabulated dataset in which an observation (row) represents an edit to a page. During the process, we collated edits from the same contributor into a single observation if they were no more than 8 hours apart, on the same page and no other contributor had edited the page in between. We assume that in such a case the edits represent intermediary save clicks within an overall editing session. Also, we measured the difference between consecutive page versions using a simple edit distance metric, which can be understood as ‘the minimal number of insertions, deletions and substitutions to make two strings equal’ (Navarro 2001, p. 37; see also Levenshtein, 1966). The value of the metric is 0 if the compared page versions are equal, otherwise a positive integer. The dataset transformation was done by a program written in Python programming language and resulted in a main dataset that contains 245,982,465 observations. A total of 233 pages that exceeded 1 gigabyte in size were excluded during the processing due to the limitations of our computing environment.

Second, the main dataset was divided into 108 monthly snapshots from January 2001 to December 2009 and summarized by individual contributors. Contributors can be identified by a unique user ID in each edit record (that does not necessarily reveal the real identity of the person) unless the edit was made anonymously. Therefore, an observation in a snapshot represents the aggregate activity of an individual contributor during a one-month period. The activity is described as the number of edits and pages edited, median edit distance and article length change, the number of reverted edits, and reverts done. We look at data from four namespaces: encyclopedia articles, article talk pages, policies and guidelines, and policies and guidelines talk pages.

Third, we used SAS statistical package version 9.13 to perform a non-parametric cluster analysis to observe the appearance of distinct contributor groupings based on editing behaviour. The idea of cluster analysis is to identify natural groupings in multidimensional data by minimizing intra-group differences and by maximizing inter-group differences between the cases. While most clustering algorithms require the researcher to decide the number of observed clusters either in advance or on the basis of an intermediary output, we wanted to know if and when distinct groupings could be seen to appear in the contributor base. For this purpose, the MODECLUS procedure in SAS allows setting a pseudo p -value for a clustering solution. The value can be interpreted as the likelihood of observed clusters appearing as the result of random variation in the data, although it is not strictly equivalent to the p -value in common statistical tests. We set the value at 0.05, which reflects a value commonly used in statistical analyses, and performed the analysis for each snapshot to observe how contributor clusters appear and evolve over time.

Appendix 2

Links to Wikipedia Pages

The theoretical narrative refers both directly and indirectly to a number of Wikipedia pages that are listed in the order of appearance below. The links point to the specific historical version of the page in the version control system.

- [1] **“Wikipedia:Policies and guidelines”** page on 1 November 2001
http://en.wikipedia.org/w/index.php?title=Wikipedia:Policies_and_guidelines&oldid=294531
- [2] **“Wikipedia:Ignore all rules”** page on 17 April 2002
http://en.wikipedia.org/w/index.php?title=Wikipedia:Ignore_all_rules&oldid=54587
- [3] **“Wikipedia:Three-revert rule”** page on 2 March 2006
http://en.wikipedia.org/w/index.php?title=Wikipedia:Three-revert_rule&oldid=41964482
- [4] **“Wikipedia:Featured articles”** page on 1 August 2006
http://en.wikipedia.org/w/index.php?title=Wikipedia:Featured_articles&oldid=66948808
- [5] **“Wikipedia:Editorial oversight and control”** page on 8 January 2007
http://en.wikipedia.org/w/index.php?title=Wikipedia:Editorial_oversight_and_control&oldid=99233111
- [6] **“Opening Plenary (transcript)”** page on 6 February 2007 from Wikimania 2006 conference wiki
[http://wikimania2006.wikimedia.org/w/index.php?title=Opening_Plenary_\(transcript\)&oldid=11012](http://wikimania2006.wikimedia.org/w/index.php?title=Opening_Plenary_(transcript)&oldid=11012)
- [7] **“User:ClueBot NG”** page on 20 October 2010
http://en.wikipedia.org/w/index.php?title=User:ClueBot_NG&oldid=391868393
- [8] **“Wikipedia:Bot policy”** page on 17 January 2007
http://en.wikipedia.org/w/index.php?title=Wikipedia:Bot_policy&oldid=101295880
- [9] **“Wikipedia:Editorial oversight and control”** page on 19 October 2007
http://en.wikipedia.org/w/index.php?title=Wikipedia:Editorial_oversight_and_control&oldid=165646130
- [10] **“Wikipedia:Flagged revisions”** page on 14 September 2013
http://en.wikipedia.org/w/index.php?title=Wikipedia:Flagged_revisions&oldid=572838257