

Coase's Penguin, or, Linux and The Nature of the Firm

Yochai Benkler*

Abstract

For decades our common understanding of the organization of economic production has been that individuals order their productive activities in one of two ways: either as employees in firms, following the directions of managers, or as individuals in markets, following price signals. This dichotomy was first identified in the early work of Ronald Coase and was developed most explicitly in the work of institutional economist Oliver Williamson. In this paper I explain why we are beginning to see the emergence of a new, third mode of production, in the digitally networked environment, a mode I call commons-based peer production.

In the past three or four years, public attention has focused on a fifteen-year-old social-economic phenomenon in the software development world. This phenomenon, called free software or open source software, involves thousands or even tens of thousands of programmers contributing to large and small scale projects, where the central organizing principle is that the software remains free of most constraints on copying and use common to proprietary materials. No one “owns” the software in the traditional sense of being able to command how it is used or developed, or to control its disposition. The result has been the emergence of a vibrant, innovative and productive collaboration, whose participants are not organized in firms and do not choose their projects in response to price signals.

This paper explains that while free software is highly visible, it is in fact only one example of a much broader social-economic phenomenon. I suggest that we are seeing the broad and deep emergence of a new, third mode of production in the

* Professor of Law, New York University School of Law. Research for this paper was partly supported by a grant from the Filomen D’Agostino and Max Greenberg Fund at NYU School of Law. I owe thanks to many for comments on this and earlier drafts, including: Bruce Ackerman, Ed Baker, Elazar Barkan, Dan Burk, Jamie Boyle, Niva Elkin Koren, Terry Fisher, Natalie Jeremijenko, Dan Kahan, Doug Lichtman, Tara Lemmy, Mark Nadel, Carol Rose, Bob Ellickson, Peggy Radin, Clay Shirky, Helen Nissenbaum, Jerry Mashaw, Eben Moglen, Larry Lessig, Chuck Sabel, Alan Schwartz, Richard Stallman, and Kenji Yoshino. I owe special thanks to Steve Snyder for his invaluable research assistance on the peer production enterprises described here.

I have gotten many question about the “Coase’s Penguin” portion of the title. It turns out that the geek culture that easily recognizes “Coase” doesn’t recognize the “penguin,” and vice versa. So, “Coase” refers to Ronald Coase, who originated the transactions costs theory of the firm that provides the methodological template for the positive analysis of peer production that I offer here. The penguin refers to the fact that the Linux kernel development community has adopted the image of a paunchy penguin as its mascot/trademark. One result of this cross-cultural conversation is that I will occasionally explain in some detail concepts that are well known in one community but not in the other.

digitally networked environment. I call this mode “commons-based peer production,” to distinguish it from the property- and contract-based modes of firms and markets. Its central characteristic is that groups of individuals successfully collaborate on large-scale projects following a diverse cluster of motivational drives and social signals, rather than either market prices or managerial commands.

I explain why this mode has systematic advantages over markets and managerial hierarchies when the object of production is information or culture, and where the physical capital necessary for that production—computers and communications capabilities—is widely distributed instead of concentrated. In particular, this mode of production is better than firms and markets for two reasons. First, it is better at identifying and assigning human capital to information and cultural production processes. In this regard, peer production has an advantage in what I call “information opportunity cost.” That is, it loses less information about who the best person for a given job might be than either of the other two organizational modes. Second, there are substantial increasing returns, in terms of allocation efficiency, to allowing larger clusters of potential contributors to interact with large clusters of information resources in search of new projects and opportunities for collaboration. Removing property and contract as the organizing principles of collaboration substantially reduces transaction costs involved in allowing these large clusters of potential contributors to review and select which resources to work on, for which projects, and with which collaborators. This results in the potential for substantial allocation gains. The article concludes with an overview of how these models use a variety of technological and social strategies to overcome the collective action problems usually solved in managerial and market-based systems by property, contract, and managerial commands.

Introduction

Imagine that back in the days when what was good for GM was good for the country an advisory committee of economists had recommended to the President of the United States that the federal government should support the efforts of volunteer communities to design and build their own cars, either for sale or for free distribution to automobile drivers. The committee members would probably have been locked up in a psychiatric ward—if Senator McCarthy or the House Un-American Activities Committee did not get them first. Yet, in September of 2000, something like this in fact happened. The President’s Information Technology Advisory Committee recommended that the federal government back open source software as a strategic national choice to sustain the U.S. lead in critical software development.¹

¹ President’s Information Technology Advisory Committee, Developing Open Source Software to Advance High End Computing, October, 2000. <http://www.ccic.gov/pubs/pitac/pres-oss-11sep00.pdf>

At the heart of the economic engine of the world's most advanced economies, and in particular that of the United States, we are beginning to take notice of a hardy, persistent, and quite amazing phenomenon—a new model of production has taken root, one that should not be there, at least according to our most widely held beliefs about economic behavior. It should not, the intuitions of the late 20th century American would say, be the case that thousands of volunteers will come together to collaborate on a complex economic project. It certainly should not be that these volunteers will beat the largest and best financed business enterprises in the world at their own game. And yet, this is precisely what is happening in the software world.

The emergence of free software,² and the phenomenal success of its flagships—the GNU/Linux operating system,³ the Apache web server, Perl, sendmail, BIND—and many others,⁴ should force us to take a second look at the dominant paradigm we hold about productivity. In the late 1930s, Ronald Coase wrote his article, *The Nature of the Firm*,⁵ in which he explained why firms—clusters of resources and agents that interact through managerial command systems rather than markets—emerge. In that paper Coase introduced the concept of transaction costs—that is, that there are costs associated with defining and enforcing property and contract rights—which are a necessary incident of organizing any activity on a market model. Coase explained the emergence and limits of firms based on the differences in the transaction costs associated with organizing production through markets or through firms. People would use the markets when the gains from doing so, net of transaction costs, exceed the gains from doing the same thing in a managed firm, net of the organization costs. Firms would emerge when the opposite was true. Any individual firm would stop growing when its organization costs exceeded the

² I use the terms “free software” and “open source software” interchangeably in this article. Those who consider the phenomenon as first and foremost involving political values, to wit, freedom, use the former, in self-conscious contradistinction to those who focus on the economic significance. See <http://www.fsf.org/philosophy/free-software-for-freedom.html>, Eric Raymond, Homesteading the Noosphere, 2 (2000) available <http://www.tuxedo.org/~esr/writings/cathedral-bazaar/homesteading/>. I have written and continue to write quite extensively of the normative implications of how information production is organized, see, e.g., Benkler, The Battle Over the Institutional Ecology of the Digitally Networked Environment, 44(2) Communications of the ACM 84 (2000), but not in this paper. I generally abjure disputations over the word.

³ I describe the operating system as GNU/Linux to denote that it is a combination of the kernel development project initiated by Linus Torvalds in 1991, and of many other operating system components created by the GNU project, originated by Richard Stallman, the father of free software, in 1984. Throughout the paper, I refer to GNU or Linux separately, to denote the specific development project, and to the operating system as GNU/Linux. I departed from this practice in the title for stylistic purposes alone. The complete GNU/Linux operating system is what everyone has in mind when they speak of the breathtaking success of free software at making excellent high-end software.

⁴ For an excellent history of the free software movement and of the open source development methodology see Glyn Moody, *Rebel Code* (2001).

⁵ Ronald H. Coase, *The Nature of the Firm*, 4 *Economica* 386 (1937).

organization costs of a newly formed, smaller firm. This basic insight was then extended and developed in the work of Oliver Williamson and other institutional economists who studied the relationship between markets and managerial hierarchies as models of organizing production.⁶

The emergence of free software as a substantial force in the software development world poses a puzzle for this conception of organization theory. Free software projects do not rely either on markets or on managerial hierarchies to organize production. Programmers do not generally participate in a project because someone who is their boss told them to, though some do. They do not generally participate in a project because someone offers them a price to do so, though some participants do focus on long-term appropriation through money-oriented activities, like consulting or service contracts. But the critical mass of participation in projects cannot be explained by the direct presence of a price or even a future monetary return, particularly in the all-important micro-level decisions regarding selection of projects to which participants contribute.⁷ In other words, programmers participate in free software projects without following the normal signals generated by market-based, firm-based, or hybrid models.

⁶ The initial framing in terms of the opposition between markets and hierarchy was Williamson's. See Oliver Williamson, *Markets And Hierarchies: Analysis and Antitrust Implications: A Study In The Economics Of Internal Organization* (1975), Oliver Williamson, *The Economic Institutions Of Capitalism* (1985), Benjamin Klein et al., *Vertical Integration, Appropriable Rents, and the Competitive Contracting Process*, 21 J.L. & Econ. 297 (1978). State hierarchies are also an option, and while the extreme version—socialist production—is largely discredited, some state production of some goods, like power, is still very much in play. Here I focus only on market production, however, whether decentralized and price-driven, or firm-based and managed. Any arguments about the importance of governmental investment in science, research, and the arts are independent of the potential conclusions for intellectual property that this paper suggests.

⁷ Even if it could be established, as it has not, that most contributors to free software development projects were motivated by extrinsic monetary motivations, like gaining consulting contracts through reputation and human capital gains, price would still be of small explanatory value if those motivations led to a general willingness to contribute to *some* project, but did not direct the actual selection of projects and type of contribution. In this regard, it is revealing that while reputation is perhaps the most readily available and widely cited extrinsic motivator to contribution, its explanatory force wanes when compared to the practices of two of the most successful free software projects. Neither the Apache project nor the Free Software Foundation publish individual contributions to the code they bless with personal attribution. It is possible that reputation creation and flow is a more complex social phenomenon within the high priesthood than would be implied by explicit attribution, or that the star status of the highest priests is enough to generate a reputation-based reward. It is also possible—indeed likely—that people's motivations are heterogeneous, and that some people are more driven by explicit reputation gains than others. Whether those will indeed cluster in projects where explicit reputation rewards are better organized remains a question that has not yet been studied empirically.

This puzzle has attracted increasing attention from economists⁸ and participants in the practice⁹ trying to understand their own success and its sustainability given widespread contrary intuitions. Lerner & Tirole present the best overarching view of the range of diverse micro-motivations that drive free software developers.¹⁰ This diversity of motivations, somewhat more formalized and generalized, plays an important role in my own analysis. Some writing by both practitioners and observers, supporters and critics, has focused on the “hacker ethic,” and analogized the sociological phenomenon to gift exchange systems.¹¹ Other writing has focused on the special characteristics of software as an object of production.¹² In this paper I approach this puzzle by departing from free software. Rather than trying to explain what is special about software or hackers, I generalize from the phenomenon of free software to suggest what makes large scale collaborations in many information production fields sustainable and productive in the digitally networked environment without reliance either on markets or on managerial hierarchy.¹³ Hence the title of the article, to invoke the challenge that the paunchy

⁸ An excellent overview of, and insightful contribution to this literature is a working paper by Steven Weber, *The Political Economy of Open Source*, (BRIE Working Paper No. 140), available <http://brie.berkeley.edu/~briewww/pubs/wp/wp140.pdf>.

⁹ The canonical references here are to two works by Eric Raymond, and open source software developer who turned into the most vocal and widely read commentator on the phenomenon, Eric Raymond, *The Cathedral and the Bazaar*, (1998), available at <http://www.tuxedo.org/~esr/writings/cathedral-bazaar/>; Eric Raymond, *Homesteading the Noosphere*, (1998), available http://www.firstmonday.dk/issues/issue3_10/raymond/.

¹⁰ Josh Lerner and Jean Tirole, *Some Simple Economics of Open Source*, 50 *J. Indus. Econ.* No. 2 (2002). Eric von Hippel in particular has provided both theoretical and empirical support for the importance of the use value gained by users in a user-driven innovation environment, both in software and elsewhere, see e.g., Eric von Hippel, *Innovation by User Communities: Learning from Open Source Software* 42 *Sloan Management Review* 82 (2001), and many collaborative papers available at <http://web.mit.edu/evhippel/www/Publications.htm>. See also Jean-Michael Dalle and Jean Nicolas, *'Libre' Software: Turning Fads Into Institutions*. (Working Paper 2001). <http://opensource.mit.edu/papers/Libre-Software.pdf>.

¹¹ In addition to Raymond, *supra* note 9, supporters of the sustainability of free software development who have used this framework include Rishab A. Ghosh, *Cooking Pot Markets: an economic model for the trade in free goods and services on the Internet*. 3 *First Monday*, no. 3. (1998), http://www.firstmonday.dk/issues/issue3_3/ghosh/index.html; Peter Kollock, *The Economies of Online Cooperation: Gift Exchange and Public Goods in Cyberspace*, in *Communities in Cyberspace*, (M. A. Smith & P. Kollock, eds. 1999). Less sanguine views of this development model, based on the same explanatory framework, include R. L. Glass, *The sociology of open source: of cults and cultures*. *IEEE Software*, (May/Jun 2000); David Lancashire, *Code, Culture and Cash: The Fading Altruism of Open Source Development*, 6 *First Monday* (2001), available http://www.firstmonday.org/issues/issue6_12/lancashire/index.html. See also Pekka Himanen with Linus Torvalds and Manuel Castells, *The Hacker Ethic* (2001) (on the hacker ethic generally, not solely in context of free software development).

¹² See James Bessen, *Open Source Software: Free Provision of Complex Public Goods* <http://www.researchoninnovation.org/opensrc.pdf> (2001).

¹³ The most closely related work in the open source software literature is the mapping of diverse motivations, see *supra*, note 8, and those papers that try to explain the open source software development

penguin that is the mascot of the Linux kernel development community poses for the view of organization rooted in Coase's work.¹⁴

Part I begins to tell the tale of the more general phenomenon through a number of detailed stories. Tens of thousands of individuals collaborate in five-minute increments to map Mars's craters, fulfilling tasks that would normally require full time PhDs. A quarter of a million people collaborate on creating the most important news and commentary site currently available on technology issues. Twenty-five thousand people collaborate to create a peer-reviewed publication of commentary on technology and culture. Forty thousand people collaborate to create a more efficient human-edited directory for the Web than Yahoo. I offer other examples as well. The point of Part I is simple. The phenomenon of large and medium scale collaborations among individuals, organized without markets or managerial hierarchies, is emerging everywhere in the information and cultural production system. The question is how we should understand these instances of socially productive behavior: how we should think about their economic value, how we should understand the dynamics that make them possible and make them tick.

My basic framework for explaining these emerging phenomena occupies Part II of the article. Collaborative production systems pose an information problem. The question that individual agents in such a system need to solve in order to be productive is what they should do. Markets solve this problem by attaching price signals to alternative courses of action. Firms solve this problem by assigning different signals, from different agents, different weight. To wit, what a manager says

model in terms of its information sharing characteristics. See Justin Pappas Johnson, Economics of Open Source Software. Working Paper (2000) <http://opensource.mit.edu/papers/johnsonopensource.pdf> (recognizing superior access to talent pool, but cautioning that free riding will lead to underutilization); Bruce Kogut, & Anca Metiu, Distributed Knowledge and the Global Organization of Software Development. Working Paper (2001) <http://opensource.mit.edu/papers/kogut1.pdf> (claiming that the value of a globally distributed skill set will loosen the grip of the richest countries on innovation).

¹⁴ The only treatment that specifically uses aspects of Coase's *Nature of the Firm* as an analytic framework for understanding free software is David McGowan, The Legal Implications of Open Source Software 2001 U. Ill. L. Rev. 241. Congruent with Coase's conclusion, McGowan assumes that in the absence of markets hierarchical control is necessary to coordinate the projects, and he demonstrates this effect as applied to the Linux kernel development process. He then analyzes how the licensing provisions and the social motivations and relationships involved in open source software projects form the basis for the hierarchical aspects of this software development model. See *id.*, 275-88. My own use of Coase's insights is very different. See *infra*, Part II. I apply Coase's insight regarding the centrality of comparative transaction costs to the organizational form a production enterprise will take. In my model, "information opportunity costs" plays a similar role in describing the comparative social cost of different organizational forms to the role played by transaction costs more generally in the Coasian framework. Peer production emerges, as firms do in Coase's analysis, because it can have lower information opportunity costs under certain technological/economic conditions. McGowan's paper therefore primarily intersects with this paper where I suggest that the integration in peer production processes sometimes takes the form of a hierarchy, a phenomenon that his paper studies in detail.

matters. In order to perform these functions, both markets and firms need to specify the object of the signal sufficiently so that property, contract, and managerial instructions can be used to differentiate between agents, efforts, resources, and potential combinations thereof. Where agents, effort, or resources cannot so be specified, they cannot be accurately priced or managed. The process of specification creates two sources of inefficiency. First, it causes some degree of information loss, as the transaction costs associated with specification of the characteristics of each individual human and material resource and opportunity for utilization limit the attainability of perfect specification. Second, property and contract make clusters of agents and resources sticky. A firm's employees will more readily work with a firm's owned resources than with other sources, and more readily collaborate with other employees of the firm than with outsiders. It is not impossible, obviously, to acquire and trade resources and collaborations, but it is done only when the perceived gains outweigh the transaction costs. It is in correcting these two failures that non-proprietary production strategies can improve on markets and firms.

Commons-based peer production, the emerging third model of production I describe here, relies on decentralized information gathering and exchange to reduce the uncertainty of participants, and has particular advantages as an information process for identifying human creativity available to work on information and cultural resources in the pursuit of projects, and as an allocation process for allocating that creative effort.¹⁵ It depends on very large aggregations of individuals independently scouring their information environment in search of opportunities to be creative in small or large increments. These individuals then self-identify for tasks and perform them for complex motivational reasons that I discuss at some length. If the problems of motivation and organization can be solved, however, then such a system has two major advantages over firms and markets. First, it places the point of decision about assigning any given person to any given set of resources with the individual. Given the high variability among individuals in terms of creativity, motivation, focus at any given point, availability etc., human creativity is an especially difficult resource to specify for efficient contracting or management. Firms recognize this, and attempt to solve this problem by creating various incentive compensation schemes and intangible reward schemes, like employee of the month awards. These schemes work to some extent to alleviate the information loss associated with managerial production, but

¹⁵ This third mode of production is in some measure similar to the artisan mode of production identified by the path breaking work of Michael Piore and Charles Sable, *The Second Industrial Divide* (1984). There are, however, sufficient qualitative differences that make this a new phenomenon that requires its own set of understandings, rather than a latter-day artisan cooperative. Most important are the scale of these collaborations, the absence of entry barriers in many or most of them, and the absence of direct appropriation of the products. With regards to organization literature, commons-based peer production stands in a similar relationship to artisan production as, in the property literature, commons relate to common property regimes. These are phenomena that share common characteristics, but ultimately diverge in central characteristics that require different explanations.

only insofar as a firm's agents and resources are indeed the best, and only insofar as these schemes capture all the motivations and contributions accurately. What peer production does is provide a framework, within which individuals who have the best information available about their own fit for a task can self-identify for the task. This provides an information gain over firms and markets, but only if the system develops some mechanism to filter out mistaken judgments agents make about themselves. This is why practically all successful peer production systems have a robust mechanism for peer review or statistical weeding out of contributions from agents who misjudge themselves.

As important as the information gains of peer production are its allocation gains, assuming that the variability in talent, wisdom, focus, and creativity among human agents makes individuals highly diverse in their fit for the job of converting various existing information and cultural resources into new projects. Human creativity cannot be assumed to be an on-off switch of suitability to a job, as simple industrial treatments of labor might have. One cannot say, as one might have in an industrial context, "this person passes threshold suitability requirements to pull this lever all day," and ignore variability beyond that fact. It is more likely that different people will be more or less productive with any given set of resources and collaborators for any given set of projects, and that this variability is large. I describe this diversity as a probability that any agent has of being a good fit with a set of resources and agents to produce highly valuable new information or cultural goods. Peer production has an advantage over firms and markets because it allows larger groups of individuals to scour larger groups of resources in search of materials, projects, collaborations, and combinations than do firms or individuals who function in markets. This is because when production is organized on a market or firm model, transaction costs associated with property and contract limit the access of people to each other, to resources and to projects, but do not do so when it is organized on a peer production model.¹⁶ Because fit of people to projects and to each other is variable, there are increasing returns to the scale of the number of people, resources, and projects capable of being combined. These returns are more than proportional, reflecting the increased probability that the best combination will emerge from larger clusters than from smaller clusters and result in a substantially higher probability that the right resource and human creative effort will be allocated to the project at which they would best be combined.

¹⁶ This is not to say that there are no transaction costs associated with peer production, which largely fall under the rubric of "integration" that I describe in Part III.2. It is merely to say that the transaction costs are of a different type. They may undermine the successful integration of a project, or may make participation too costly for contributors, but they do not arise as a barrier to prevent many individuals from collaborating in the same resource space, or many resources from populating that space.

The advantages of peer production, then, are improved information about, and allocation of, human creativity. These advantages can become salient, as they appear to have become, when that component of production itself becomes salient. In the domain of information and culture, production generally comprises the combination of pre-existing information/cultural inputs, human creativity, and the physical capital necessary to fix ideas and human utterances in media capable of storing and communicating them, and in transmitting them. Existing information and culture are a public good in the strict economic sense of being non-rival.¹⁷ The cost of physical capital was, for over 150 years, the central organizing principle of information and cultural production, from the introduction of high-cost, high-volume mechanical presses, through telegraph, telephone, radio, television, cable, and satellite systems. These costs largely structured production around a capital-intensive, industrial model. The declining price of computation, however, has inverted the capital structure of information and cultural production. Desktop PCs and the digital video and audio systems that connect to them are capable of performing most of the physical capital functions that once required high investments. In this context, where physical capital both for fixation and communication is low and widely distributed, and where the primary non-human input—existing information—is itself a public good, the primary remaining scarce resource is human creativity. And it is under these conditions that the relative advantages of peer production in organizing that input permit that mode of organization to emerge to much greater glory than possible before.

Obviously, this leaves the motivation and organization questions. These generally would fall under the “tragedy of the commons” critique, which I purposefully invoke by calling the phenomenon “commons-based” peer production. The traditional objections to the commons are primarily two. First, no one will invest in a project if they cannot appropriate its benefits. That is, motivation will lack. Second, no one has the power to organize collaboration in the use of the resource. That is, organization will lack and collaboration will fail. The past decade or so, however, has seen an important emerging literature on some successful commons, and

¹⁷ While the reference to information as a public good is common, the reference to culture is not. I have no intention, or need, to go into subtle definitions of culture here, though I tend to follow the approach offered in Jack Balkin, *Cultural Software* (2001), and to think of culture as a framework for comprehension. By “culture” I mean a set of representations, conceptions, interpretations, knowledge of social behavior patterns, etc, whose particular application to reducing uncertainty for human action is too remote to be called “information,” but which is indispensable to the way we make sense of the world. “Cultural production” as I use it here is what parents, and teachers, Hollywood, Mozart, or the Pope, peer groups and the guys playing guitars in Washington Square Park do. Defined as a set of conceptions and their representations, and of behavioral instructions sets, its economic character as similar to ideas or information is fairly obvious. Obviously, embodiments of culture like a specific statue or building are no more non-rival than embodiments of any other form of information, like a book or a corkscrew.

mostly on successful common property regimes.¹⁸ These involve primarily the introduction of a variety of non-property-based schemes for structuring cooperation among relatively limited groups of participants. While useful as insights in many ways into how formal and informal norms can structure collaboration, these studies of common appropriation regimes do not give a complete answer to the sustainability of motivation and organization for the truly open, large scale non-proprietary peer production projects we see on the Net.

My answer to these problems occupies Part III. The motivation problem is solved by two distinct analytic moves. The first involves the proposition that diverse motivations move human beings, and more importantly, that there exist ranges of human experience in which the presence of monetary rewards is inversely related to the presence of other, social-psychological rewards. Money, love, and sex offer an obvious and stark example, but the tradeoffs that academics face between selling consulting services, on the one hand, and writing within a research agenda respected by peers, on the other hand, are also reasonably intuitive. Given these propositions, it becomes relatively straightforward to see that there will be conditions under which a project that can organize itself to offer social-psychological rewards removed from monetary rewards will attract certain people, or certain chunks of people's day, that monetary rewards would not. The particular characteristics of large-scale collaborations that generate these conditions predictably and ubiquitously are the second component of the analysis.

The second analytic move involves understanding that when a project of any size is broken up into little pieces, each of which could be performed by an individual in a short amount of time, the motivation necessary to get any given individual to contribute need only be very small. This suggests that peer production will thrive where projects have three characteristics. First, they must be modular. That is, they must be divisible into components, or modules, each of which can be produced independently of the production of the others. This enables production to be incremental and asynchronous, pooling the efforts of different people, with different capabilities, who are available at different times. Second, the granularity of the modules is important. Granularity refers to the sizes of the project's modules, and in order for a peer production process successfully to pool a relatively large pool of contributors the modules should be predominately fine-grained, or small in size. This allows the project to capture contributions from large numbers of contributors whose motivation level will not sustain anything more than quite small efforts towards the project. Novels, for example, at least those that look like our current conception of a

¹⁸ For discussions of commons see, Carol Rose, *The Comedy of the Commons: Custom, Commerce, and Inherently Public Property*, 53 U. Chi. L. Rev. 711 (1986); Elinor Ostrom, *Governing the Commons* (1992). A brief discussion of these concepts as applied to peer production follows below, pages 35-36.

novel, are likely to prove resistant to peer production.¹⁹ In addition, a project will likely be more efficient if it can accommodate variously sized contributions. Heterogeneous granularity will allow people with different levels of motivation to collaborate by contributing smaller or larger grained contributions, consistent with their level of motivation. Third, and finally, a successful peer production enterprise must have low-cost integration, which includes both quality control over the modules and a mechanism for integrating the contributions into the finished product. If a project cannot defend itself from incompetent or malicious contributions and integrate the competent modules into a finished product at sufficiently low cost, integration will either fail or force the integrator to appropriate the residual value of the common project—usually leading to a dissipation of the motivations to contribute *ex ante*. Automated integration and iterative peer production of integration, for example the use of free software to integrate peer production of some other information good, are the primary mechanisms by which peer production projects described in this paper have lowered the cost of integration to the point where they can succeed and sustain themselves. As for a project's mechanisms for defending itself from incompetent or malicious contributions, one sees peer production enterprises using a variety of approaches towards solving collective action problems that are relatively familiar from the commons literature offline. These include various formal rules, like the GNU GPL²⁰ that prevents defection²¹ from many free software projects, including most prominently the flagship, GNU/Linux. They also include technical constraints

¹⁹ The most successful novel-like enterprise on the Net that I know of is "The Company Therapist," <http://www.tetherapist.com/>. There, the collaborative fiction problem was solved by building a system that enabled anyone to contribute a small chunk—patient's interview notes, therapist comments, etc.—to the company therapist's files, and the common project is to create a fascinating mosaic of people and stories seen through the eyes of a company therapist. Most collaborative fiction sites, however, suffer from the fact that modularity and granularity lead to disjunction relative to our expectations from novels.

²⁰ The GNU General Public License, <http://www.gnu.org/licenses/licenses.html#TOCGPL> is the most important institutional innovation of the Free Software Foundation founded by Richard Stallman. It provides the institutional constraint that prevents defection from free software projects that would take the form of taking code others have written and combining it with one's own code and then releasing it under license terms that do not allow users the same freedoms to use as the original free software did. This license does not prevent commercial distribution of free software for a fee. It places certain limits on how the software can be used as an input into derivative works that would be made less free than the original. In this, it radically breaks from the concept of the public domain that underlies copyright's law general background rule for non-proprietary materials. For discussions of the GPL, its legal nature and institutional characteristics see, Eben Moglen, Free Software Matters I, <http://moglen.law.columbia.edu/publications/lu-12.html>, and Eben Moglen, Free Software Matters II, <http://moglen.law.columbia.edu/publications/lu-13.html>. Moglen's views are particularly important since he has been General Counsel to the Free Software Foundation for the past decade and has more experience with enforcing this license than anyone else. More detailed academic treatments include McGowan, *supra*, note __; Margaret Jane Radin & R. Polk Wagner, The Myth of Private Ordering: Rediscovering Legal Realism in Cyberspace, 73 Chi.-Kent L. Rev. 1295 (1998).

²¹ I use the term "defection" to describe any action that an agent who participates in a cooperative enterprise can take to increase his or her own benefit from the common effort in a way that undermines the success or integrity of the common effort.

that prevent or limit the effect of defection. Social norms too play a role in sustaining some of these collaborations, both where there are small groups, and where there are larger groups and the platform allows for good monitoring and repair when individuals defect. Finally, the sheer size of some of these projects enables the collaboration platform to correct for defection by using redundancy of contributions and averaging out of outliers—be they defectors or incompetents.

The normative implications of recognizing peer production are substantial. At the level of political morality, what is at stake is the shape of freedom and equality in the emerging social-technological condition we associate with the Internet. These can take radical forms, both anarchistic and libertarian, as they do in the work of Eben Moglen, who was first to identify the phenomenon I now call peer production in *Anarchism Triumphant*,²² and in the minds of many in the free software community.²³ But the stakes for freedom and equality are high for a wide range of liberal commitments.²⁴ At the level of institutional design, the emergence of commons-based peer production adds a new and deep challenge to the prevailing policy of rapid expansion of the scope of exclusive rights in information and culture that has been the predominant approach in the past 25 years, as Boyle's work on the second enclosure movement elegantly elucidates,²⁵ and the dynamic of decentralized innovation plays a central role in Lessig's forceful argument for embedding the openness of commons in the architecture of the Net.²⁶ In this article, however, I do not attempt to add to the normative literature. Instead, the article is intended as a purely descriptive account of the scope of the empirical phenomenon and its analytic drivers.

One important caveat is necessary. I am *not* suggesting that peer production will supplant markets or firms. I am not suggesting that it is always the more efficient model of production for information and culture. What I am saying is that this

²² Eben Moglen, *Anarchism Triumphant* (1999)

<http://emoglen.law.columbia.edu/publications/anarchism.html>. The descriptive insight in that paper that corresponds to peer production is the phenomenon he calls there Moglen's Metaphorical Corollary to Faraday's Law:

"Moglen's Metaphorical Corollary to Faraday's Law says that if you wrap the Internet around every person on the planet and spin the planet, software flows in the network. It's an emergent property of connected human minds that they create things for one another's pleasure and to conquer their uneasy sense of being too alone."

²³ Canonical is, of course, Richard M. Stallman, *Philosophy of the GNU Project*, <http://www.gnu.org/philosophy/philosophy.html>.

²⁴ I outline the breadth of the range of liberal convictions affected by these issues in *Freedom in the Commons: Towards A Political Economy of Information*, *forthcoming* Duke L. J. (2003).

²⁵ James Boyle, *The Second Enclosure Movement and the Construction of the Public Domain*, (paper for the "Conference on the Public Domain," Duke Law School, Durham, North Carolina, November 9-11, 2001).

²⁶ Lawrence Lessig, *The Future of Ideas: The Fate of the Commons in a Connected World* (Random House 2001).

emerging third model is (a) distinct from the other two, and (b) has certain systematic advantages over the other two in identifying and allocating human capital/creativity. When these advantages will outweigh the advantages that the other two models may have in triggering or directing human behavior with the relatively reliable and reasonably well-understood triggers of money and hierarchy is a matter for more detailed study. I offer some lines of understanding the limitations of this model of production in Part III, but do not attempt a full answer to these questions here.

I. Peer Production All Around

While open source software development has captured the attention and devotion of many, it is by no stretch of the imagination the first or most important instance of production by peers who interact and collaborate without being organized on either a market-based or a managerial/hierarchical model. Most important in this regard is the academic enterprise, and in particular scientific research. Thousands of individuals make individual contributions to a body of knowledge, set up internal systems of quality control, and produce the core of our information and knowledge environment. These individuals do not expect to exclude from their product anyone who does not pay for it, and for many of them the opportunity cost of participating in academic research, rather than applying themselves to commercial enterprise, carries a high economic price tag. In other words, individuals produce on a non-proprietary basis, and contribute their product to a knowledge “commons” that no one is understood as “owning,” and that anyone can, indeed is required by professional norms to, take and extend. We appropriate the value of our contributions using a variety of service-based rather than product-based models (teaching rather than book royalties) and grant funding from government and non-profit sources, as well as, at least as importantly, reputation and similar intangible—but immensely powerful—motivations embodied in prizes, titles etc. It is easy, though unjustifiable, in the excitement of moment that feels like one of great transformation to forget that information production is one area where we have always had a mixed system of commercial/proprietary and non-proprietary peer production—not as a second best or a contingent remainder from the middle ages, but because at some things the non-proprietary peer production system of the academic world is simply better.²⁷

²⁷ An early version of this position is Richard R. Nelson, *The Simple Economics of Basic Scientific Research*, 48 *Journal of Political Economy* 297-306 (June 1959); more recently one sees the work, for example, of Rebecca S. Eisenberg, *Public Research and Private Development: Patents and Technology Transfer In Government-Sponsored Research*, 82 *Va. L. Rev.* 1663, 1715-24 (1996). For a historical description of the role of market and non-market institutions in science see Paul A. David, *From Market Magic to Calypso Science Policy* (1997) (Stanford University Center for Economic Policy Research Pub. No. 485).

In one thing, however, academic peer production and commercial production are similar. Both are composed of people who are professional information producers. The individuals involved in production have to keep body and soul together from information production. However low the academic salary is, it must still be enough to permit one to devote most of one's energies to academic work. The differences reside in the modes of appropriation and in the modes of organization—in particular how projects are identified and how individual effort is allocated to projects. Academics select their own projects, and contribute their work to a common pool that eventually comprises our knowledge of a subject matter, while non-academic producers will often be given their marching orders by managers, who themselves take their focus from market studies, and the product is then sold into the market for which it was produced.

Alongside the professional model, it is also important to recognize that we have always had nonprofessional information and cultural production on a non-proprietary model. Individuals talking to each other are creating information goods, sometimes in the form of what we might call entertainment, and sometimes as a means for news distribution or commentary. Nonprofessional production has been immensely important in terms of each individual's information environment. If one considers how much of the universe of communications one receives in a day comes from other individuals in one-to-one or small-scale interactions—such as email, lunch, or hallway conversations—the effect becomes tangible.

Nonetheless, ubiquitous computer communications networks are bringing about a dramatic change in the scope, scale, and efficacy of peer production. As computers become cheaper and as network connections become faster, cheaper, and more ubiquitous, we are seeing the phenomenon of nonprofessional peer production of information scale to much larger sizes, performing more complex tasks than were in the past possible for, at least, nonprofessional production. To make this phenomenon more tangible, I will describe in this part a number of such enterprises, organized so as to demonstrate the feasibility of this approach throughout the information production and exchange chain. While it is possible to break an act of communication into finer-grained subcomponents,²⁸ largely we see three distinct functions involved in the process. First, there is an initial utterance of a humanly meaningful statement. Writing an article or drawing a picture, whether done by a professional or an amateur, whether high quality or low, is such an action. Second there is a separate function of mapping the initial utterances on a knowledge map. In particular, an utterance must be understood as “relevant” in some sense and “credible.” “Relevant” is a subjective question of mapping an utterance on the

²⁸ See Yochai Benkler, Communications infrastructure regulation and the distribution of control over content, 22 *Telecomms Policy* 183 (1998).

conceptual map of a given user seeking information for a particular purpose defined by that individual. If I am interested in finding out about the political situation in Macedonia, a news report from Macedonia or Albania is relevant, even if sloppy, while a Disney cartoon is not, even if highly professionally rendered. Credibility is a question of quality by some objective measure that the individual adopts as appropriate for purposes of evaluating a given utterance. Again, the news report may be sloppy and not credible, while the Disney cartoon may be highly accredited *as a cartoon*. The distinction between the two is somewhat artificial, however, because very often the utility of a piece of information will depend on a combined valuation of its credibility and relevance. A New York Times story on the Balkans in general, for example, will likely be preferable to excited gossip in the cafeteria specifically about Macedonia. I will therefore refer to “relevance/accreditation” as a single function for purposes of this discussion, keeping in mind that the two are complementary and not entirely separable functions that an individual requires as part of being able to use utterances that others have uttered in putting together the user’s understanding of the world. Finally, there is the function of distribution, or how one takes an utterance produced by one person and distributes it to other people who find it credible and relevant. In the mass media world, these functions were often, though by no means always, integrated. NBC news produced the utterances, gave them credibility by clearing them on the evening news, and distributed them simultaneously. What the Net is permitting is much greater disaggregation of these functions, and so this part will proceed to describe how each component of this information production chain is being produced on a peer-based model on the Net for certain information and cultural goods other than software.

1. Content

NASA Clickworkers is “an experiment to see if public volunteers, each working for a few minutes here and there can do some routine science analysis that would normally be done by a scientist or graduate student working for months on end.”²⁹ Users can mark craters on maps of Mars, classify craters that have already been marked or search the Mars landscape for “honeycomb” terrain. The project is “a pilot study with limited funding, run part-time by one software engineer, with occasional input from two scientists.”³⁰ In its first six months of operation over 85,000 users visited the site, with many contributing to the effort, making over 1.9 million entries (including redundant entries of the same craters, used to average out errors.) An analysis of the quality of markings showed “that the automatically-computed consensus of a large number of clickworkers is virtually indistinguishable from the inputs of a geologist with years of experience in identifying Mars craters.”³¹

²⁹ <http://clickworkers.arc.nasa.gov/top>

³⁰ <http://clickworkers.arc.nasa.gov/contact>

³¹ Clickworkers Results: Crater Marking Activity, July 3, 2001,

The tasks performed by clickworkers (like marking craters) are discrete, and each iteration is easily performed in a matter of minutes. As a result users can choose to work for a few minutes doing one iteration or for hours by doing many, with an early study of the project suggesting that some clickworkers indeed work on the project for weeks, but that 37% of the work was done by one-time contributors.³²

The clickworkers project is a particularly crisp example of how complex professional tasks that required budgeting the full time salaries of a number of highly trained individuals can be reorganized so as to be performed by tens of thousands of volunteers in increments so minute that the tasks can now be performed on a much lower budget. This low budget is devoted to coordinating the volunteer effort, but the raw human capital needed is contributed for the fun of it. The professionalism of the original scientists is replaced by a combination of very high modularization of the task, coupled with redundancy and automated averaging out of both errors and purposeful defections (*e.g.*, purposefully erroneous markings).³³ What the NASA scientists running this experiment had tapped in to was a vast pool of five-minute increments of human judgment applied with motivation to a task that is unrelated to keeping together the bodies and souls of the agents.

While clickworkers is a distinct, self-conscious experiment, it suggests characteristics of distributed production that are, in fact, quite widely observable. Consider, for example, how the networked environment has enabled new ways of fulfilling the function that traditionally was fulfilled by encyclopedias or almanacs. At the most general level, consider the World Wide Web itself. Individuals put up web sites with all manner of information, in all kinds of quality and focus, for reasons that have nothing to do with external, well-defined economic motives—just like the individuals who identify craters on Mars. A user interested in information need only plug a search request into a search engine like Google, and dozens, or hundreds of websites will appear. Now, there is a question of how to select among them—the question of relevance and accreditation. But that is for the next subpart. For now what is important to recognize is that the web is a global library produced by millions of people. Whenever I sit down to search for information, there is a very high likelihood that someone, somewhere, has produced a usable answer, for whatever reason—pleasure, self-advertising, or fulfilling some other public or private goal as a non-profit or for profit that sustains itself by means other than selling the information I need. The power of the web to answer such an encyclopedic question comes not from the fact that one particular site has all the great answers. It is not an

<http://clickworkers.arc.nasa.gov/documents/crater-marking.pdf>

³² B. Kanefsky, N.G. Barlow, and V.C. Gulick, Can Distributed Volunteers Accomplish Massive Data Analysis Tasks? <http://clickworkers.arc.nasa.gov/documents/abstract.pdf>

³³ Clickworkers results, *supra*, para. 2.2. (describing, among other things, the exclusion of the markings of 56 students in an art class who marked concentric circles instead of trying to mark craters).

Encyclopedia Britannica. The power comes from the fact that it allows a user looking for specific information at a given time to collect answers from a sufficiently large number of contributions. The task of sifting and accrediting falls to the user, motivated by the need to find an answer to the question posed. As long as there are tools to lower the cost of *that* task to a level acceptable to the user, the Web shall have “produced” the information content the user was looking for. These are not trivial considerations. But they are also not intractable. As we shall see, some of the solutions can themselves be peer produced, but some solutions are emerging as a function of the speed of computation and communication, which enables more efficient technological solutions.

But, one might argue, that is still not an encyclopedia, in the sense of a coherently ordered locus of a wide range of human knowledge in relatively accessible and digested form. Can that task, which requires more disciplined writing for that purpose, be performed on a distributed model? The beginning of an answer is provided by the Wikipedia project.³⁴ The project involves some two thousand volunteers who are collaborating to write an encyclopedia. The project runs on a free software collaborative authorship tool, Wiki, which is a markup language similar in concept to HTML, but relatively easier to implement and focused on permitting multiple people to edit a single document and interlocking documents, while generating archives of changes made to each. While 2000 people have not been able to generate a complete encyclopedia in roughly 18 months of operation, they have made substantial progress, producing about 30,000 articles, and readers are invited to test their own evaluation of the quality. The terms “chimpanzee,” “computational complexity theory,” or simply “copyright,” for example, provide good demonstrations. A comparison to www.encyclopedia.com, the online version of the Columbia Encyclopedia, would suggest that Wikipedia cannot yet be said to be either systematically better or worse. Given that it is a volunteer effort, and the comparison is to an established commercial encyclopedia, that is actually saying quite a bit. Perhaps the most interesting characteristic about Wikipedia is the self-conscious social-norms-based dedication to objective writing. The following fragments from the self-described essential characteristics and basic policies of Wikipedia are illustrative:

First and foremost, the Wikipedia project is self-consciously an encyclopedia - rather than a dictionary, discussion forum, web portal, etc. See “encyclopedia” as well as “what Wikipedia is not”. . . .

Wikipedia's participants commonly follow, and enforce, a few basic policies that seem essential to keeping the project running smoothly and productively. The following are just a few of those policies; for more information, please see “Wikipedia policy”.

³⁴ <http://www.wikipedia.com>

First, because we have a huge variety of participants of all ideologies, and from around the world, Wikipedia is committed to making its articles as unbiased as possible. The aim is not to write articles from a single *objective* point of view—this is a common misunderstanding of the policy—but rather, to fairly and sympathetically present all views on an issue. See “neutral point of view” page for further explanation, and for a very lengthy discussion.³⁵

The point to see from this quote is that the participants of Wikipedia are plainly people who like to write. Some of them participate in other collaborative projects, like Everything2.com.³⁶ But when they enter the common project of “Wikipedia,” they undertake to participate in a particular way—a way that the group as a group has adopted to make its product be an “encyclopedia.” On their interpretation, that means conveying in brief terms the state of the art on the item, including divergent opinions about it, but not the author’s opinion. Whether that is an attainable goal is a subject of interpretive theory, and is a question as applicable to a professional encyclopedia as it is to Wikipedia. My point, however, is that Wikipedia provides a rich example of a medium sized collection of individuals, who collaborate to produce an information product of mid- to highbrow quality, and is reasonably successful. In particular, it suggests that even in a group of this size, social norms coupled with a simple facility to allow any participant to edit out blatant opinion written in by another in contravention of the social norms keep the group on track.

³⁵ <http://www.wikipedia.com/wiki/Wikipedia>. The “neutral point of view” page is indeed revealing of how explicit and central to the project the social norm of objective-style writing is. See <http://www.wikipedia.com/wiki/wikipedia%3Aneutral+point+of+view>.

³⁶ <http://www.everything2.com>. Everything2.com is a “complex online community with a focus to write, publish and edit a quality database of information, insight and humor.” <http://everything2.com/index.pl?node=Everything%20FAQ>. The system enables registered users to post “write-ups” and create “nodes” pertaining to particular topics that they define. It does not have a directory structure, instead nodes are linked together with hypertext within the text of the node and also with a matrix of related links at the bottom of each node. The linking is done initially by the author—thereby self-generating an conceptual map, and later by others. A node is a particular topic identified by the title of the node. After the author of the first write-up creates a “nodeshell”, other users can add additional write-ups to that node. Write-ups are constantly being reviewed and removed by editors. Editors are chosen based on “merit, seniority and writing skill.” http://www.everything2.com/index.pl?node=The%20Power%20Structure%20of%20Everything%20&lastrnode_id=596635 Everything2 also contains a voting system for non-editor users to vote on each other’s write-ups. Although each write-up has a reputation based on whether it has been voted up or down, the write-up does not get automatically filtered due to a low reputation. In other words, the system combines individually authored materials and individually defined mappings of relevance of materials, with common procedures, some purely democratic, some based on a rotating hierarchy of editors appointed by experience and reputation within the system, reputation built from the collective judgments of their peers. The result is a substantial database of writings on a wide variety of topics.

Perhaps the most sophisticated locus of peer reviewed, mid- to high-quality essays published on the Net as of early 2002 is Kuro5hin, also known as K5.³⁷ Self-described as

Kuro5hin.org is a community of people who like to think. You will not find garbage in the discussions here, because noise is not tolerated. This is a site for people who want to discuss the world they live in. It's a site for people who are on the ground in the modern world, and who sometimes look around and wonder what they have wrought.³⁸

As of March 2002, it appeared that Kuro5hin had roughly 25,000 users.³⁹ Articles run a broad gamut of topics, but are supposed to be roughly around technology and culture. The general headings are Technology, Culture, Freedom & Politics, Media, News, Op-ed., Columns, Meta (dedicated to discussion of K5 itself), and MLP (mindless link propagation, a general catchall category of things the community members find interesting). The articles include news reportage from other sources, but most of the interesting materials provide some form of commentary as well. The articles and responses to them are fairly substantial.

The site and community have a fairly heavy emphasis on quality of materials published. The guide to articles' submission⁴⁰ emphasizes quality of information and writing multiple times, and prepares new contributors for the experience of quite close peer review of their submission. Beyond the general guidelines, the software that runs Kuro5hin, Scoop, a free software project initiated by one of the co-founders of K5, implements a series of steps both before and after submission and publication of an article that serve as collaborative quality control mechanisms. The emphasis on quality is enforced by the site's mechanism for peer-review pre-publication and peer-commentary post-publication.⁴¹

³⁷ <http://www.kuro5hin.org/>. The discussion here is deeply indebted to the work of Caio Mario de Silva, Kuro5hin.org, Collaborative Media, and Political Economy of Information (unpub. man. on file with author). Another source is Everett Teach, Clay Hackett, Bobby Nall, *Ethnography of Kuro5hin*, <http://ccwf.cc.utexas.edu/~hackett/k5/>.

³⁸ *Id.*, Mission Statement.

³⁹ <http://www.kuro5hin.org/story/2002/3/16/51221/8976>.

⁴⁰ <http://www.kuro5hin.org/?op=special;page=article>

⁴¹ When an article is submitted it is not automatically placed in a publicly viewable space. It is placed, instead, in a submission queue. At that point, all registered users of K5 have an opportunity to comment on the article, provide suggestions for correction and improvement, and vote their opinion whether they think the story should be placed on the front page, a specialty page, or rejected. The system determines some critical number of votes necessary for any one of these actions, based on the number of users then registered. Typically rejection requires fewer votes than acceptance. Articles may be resubmitted after being rejected, typically after having been revised in accordance with the comments. The system up to this point is remarkably similar to academic peer review in many respects, except the scope of participation and the egalitarian and democratic structure of the editorial decision. After publication, K5

A very different type of trend to look at in regard to collaborative creation is the emergence and rise of computer games, and in particular multi-player and online games. These fall in the same cultural “time slot” as television shows and movies of the 20th century. The interesting thing about them is that they are structurally different. In a game like Ultima Online or EverQuest, the role of the commercial provider is not to tell a finished, highly polished story to be consumed start to finish by passive consumers. Rather, the role of the game provider is to build tools with which users collaborate to tell a story. There have been observations about this approach for years, regarding MUDs and MOOs.⁴² The point here is that there is a

provides the platform for readers to comment on articles, and for other readers to rate these comments for their relevance and quality. The system is different in various respects from the Slashdot system described in detail in Part I.2., but the principle is the same. It permits readers to post comments. It permits other readers to rate comments as better or worse. It aggregates these individual ratings into collective judgments about the quality of comments, judgments that can then be used by the site's readers to filter out lower quality comments. In general, all the characteristics described in this section go to questions of how one generates relevance and accreditation on a peer production model, and will be explored in greater detail, in the context of other sites, in the next subsection. The point to take away at this point is that part of what makes K5 so successful in maintaining quality is a rather elaborate, large scale peer review system and post-publication commentary, which itself is then peer reviewed in an iterative process.

⁴² MUDs (Multi-user Dungeon or Multi-user Dimension) and MOOs (MUD, Object Oriented) are acronyms for software programs that create an interactive multi-user networked text-based virtual world. The software maintains a database of users and objects that the users can interact with in a variety of ways. MUDs are typically built around a theme. MUD “worlds” are often based on books, movies, cartoons and other role playing games. <http://www.geocities.com/TimesSquare/9944/>. Pavel Curtis, creator of perhaps the most famous of MOOs. LambdaMOO, identified three elements that distinguish MUDs from typical role-playing games:

A MUD is not goal-oriented; it has no beginning or end, no ‘score’, and no notion of ‘winning’ or ‘success’. In short, even though users of MUDs are commonly called players, a MUD isn't really a game at all.

A MUD is extensible from within; a user can add new objects to the database such as rooms, exits, ‘things’, and notes. Certain MUDs, including the one I run, even support an embedded programming language in which a user can describe whole new kinds of behavior for the objects they create.

A MUD generally has more than one user connected at a time. All of the connected users are browsing and manipulating the same database and can encounter the new objects created by others. The multiple users on a MUD can communicate with each other in real time.”

Pavel Curtis, Social Phenomena in Text-Based Virtual Realities, in Proceedings of the 1992 Conference on the Directions and Implications of Advanced Computing, available

<http://citeseer.nj.nec.com/curtis92mudding.html>. There are various acronyms for various types of MUD-like variations including MUD, MUSH, MUX and MUCK. All of the variations run basically the same software, the primary difference between them is how much freedom the characters have to modify the environment. All M*s (refers to all MUD-like variants) are administered in some way by those that setup the software and maintain the connectivity. Typically, the administrator will setup the initial world and implement some coded commands. The administrator will also setup a hierarchy of user level and as users advance they have more control over the objects within the game and they can create coded commands. It is these decisions - how much of the world does the administrator create, how rich are the

discrete element of the “content” that can be produced in a centralized professional manner—the screenwriter here replaces the scientist in the NASA clickworkers example—that can also be organized using the appropriate software platform to allow the story to be written by the many users as they experience it. The individual contributions of the users/co-authors of the storyline are literally done for fun—they are playing a game—but they are spending real economic goods—their attention and substantial subscriptions fees—on a form of entertainment that displaces what used to be passive reception of a finished, commercially and professionally manufactured good with a platform for active co-production of a storyline. The individual contributions are much more substantial than the time needed to mark craters, but then the contributors are having a whole lot more fun manipulating the intrigues of their imaginary Guild than poring over digitized images of faint craters on Mars.

2. Relevance/accreditation

Perhaps, you might say, many distributed individuals can produce content, but it is gobbledygook. Who in their right mind wants to get answers to legal questions

coded commands, does the administrator allow users to have the power to manipulate the game – that distinguish the various M*’s from each other. MUDs are typically heavy on coded commands and designed to be battle heavy. MUSHs on the other hand “are unlikely to have coded commands to the same extent that a MUD will, relying instead on arbitration or consent to determine the effects of actions.” MOOs are perhaps the exception in that most of them are not role-playing, but “educational and social.” <http://wso.williams.edu:8000/~msulliva/mushes/explan.html>.

Most important in the history of MUDs was LambdaMOO. “LambdaMOO is a MOO: a MUD that uses an object-oriented programming language to manipulate objects in the virtual world.”

<http://cobot.research.att.com/lambdaMOO.html>. LambdaMOO was created in 1990 by Pavel Curtis as a social experiment. “[It] is the first, most diverse, oldest, largest, and most well-known MOO.” “When Pavel Curtis took on the project of developing the MOO environment, he gave it a social focus instead of the game goal of traditional MUDs” <http://world.std.com/~rs/inevitable.html>. The original site has remained active for over a decade and continues to thrive with over 100,000 people having participated in this one virtual world. “LambdaMOO is thus a long-standing, ongoing experiment in collective programming and creation, with often stunning results that can only be fully appreciated firsthand. Inventions include technical objects, such as the lag meter, which provides recent statistics on server load; objects serving a mix of practical and metaphorical purposes, such as elevators that move users between floors; objects with social uses, such as the birthday meter, where users register their birthdays publicly; and objects that just entertain or annoy, such as the Cockatoo, a virtual bird who occasionally repeats an utterance recently overheard.” Rebecca Spainhower details the evolution of the social structure of LambdaMOO in her article *Real Problems in Virtual Communities*.

<http://world.std.com/~rs/inevitable.html>. Generally, the MOO was administered by a few system administrators (called Wizards within the game). Haakon (Pavel Curtis’ Wizard character) drafted a set of guidelines for behavior. When administration became too overwhelming for the Wizards, they appointed an “Architecture Review Board” of 15 trusted users to allocate space to new users. The Wizards were still responsible for dealing with unruly users and mediating disputes. In 1993, the Wizards turned that responsibility over to the community at large by implementing a democratic petitioning and balloting system. Since that time the community has addressed problems of population growth, harassment and the behavior of anonymous guest accounts.

from a fifteen-year-old child who learned the answers from watching Court TV?⁴³ The question, then, becomes whether relevance and accreditation of initial utterances of information can itself be produced on a peer production model. At an initial intuitive level the answer is provided by commercial businesses that are breaking off precisely the “accreditation and relevance” piece of their product for peer production. Amazon is a perfect example.

Amazon uses a mix of mechanisms to get in front of their buyers books and other products that the users are likely to buy.⁴⁴ A number of these mechanisms produce relevance and accreditation by harnessing the users themselves. At the simplest level, the recommendation “customers who bought items you recently viewed also bought these items,” is a mechanical means of extracting judgments of relevance and accreditation from the collective actions of many individuals who produce the datum of relevance as a by-product of making their own purchasing decisions. At a more self-conscious level (self-conscious, that is, on the part of the user), Amazon allows users to create topical lists, and to track other users as their “friends and favorites,” whose decisions they have learned to trust. Amazon also provides users with the ability to rate books they buy, generating a peer-produced rating by averaging the ratings. The point to take home from Amazon is that a corporation that has done immensely well at acquiring and retaining customers harnesses peer production to provide one of its salient values—its ability to allow users to find things they want quickly and efficiently.

⁴³ NY Times Magazine, Sunday, July 15, 2001 cover story.

⁴⁴ These include both automatically generated and human made relevance maps. For example, “Page you made” is based on the user’s recent clicks on the site, lists a “featured item” as well as several “quick picks” which are products that are similar to the recently viewed items. The page also contains a section called “customers who bought items you recently viewed also bought these items.” The page features “listmania” lists, which are user-created topical lists, and a “more to explore” section that provides relevant links to a topical directory of the Amazon inventory. Another mechanism that Amazon uses are Friends and Favorites, providing users the ability to track their “favorite people” (who permit themselves to be so tracked by others) and track the various product reviews that the person gives. Users can also “share purchases” and make their purchases available for other users to see. If the user finds a person with similar tastes, these options could aid with relevance and if the user finds a particularly trustworthy person it could aid in accreditation of the product. Amazon also provides discussion boards for direct exchange between users. Amazon also creates “purchase circles” which are “highly specialized bestseller lists,” based on aggregated data consisting of purchases sent to particular zip codes or ordered from a particular Internet domain name. The data is analyzed and compared to site-wide trends to come up with lists of items that are more popular with that particular group than with the general population. The groups are divided either geographically (by town or city) or organizationally (schools, government offices, corporations). If users find a list particularly useful, they can bookmark the list to view the changes as the list is updated to reflect new sales data. Recommendations: Amazon software recommends certain products to the user. These “recommendations” are based on items the user has purchased or rated, and their activity on the site contrasted with other users’ activity. As a result, the recommendations can change when the user purchases or reviews an item, or when the interests of other consumers change.

Similarly, Google, widely recognized as the most efficient general search engine currently operating, introduced a crucial innovation into ranking results that made it substantially better than any of its competitors. While Google uses a text-based algorithm to retrieve a given universe of web pages initially, its PageRank software employs peer production of ranking in the following way.⁴⁵ The engine treats links from other web sites pointing to a given website as votes of confidence. Whenever one person who has a page links to another page, that person has stated quite explicitly that the linked page is worth a visit. Google's search engine counts these links as distributed votes of confidence in the quality of that page as among pages that fit the basic search algorithm. Pages that themselves are heavily linked-to count as more important votes of confidence, so if a highly linked-to site links to a given page, that vote counts for more than if an obscure site that no one else thinks is worth visiting links to it. The point here is that what Google did was precisely to harness the distributed judgments of many users, each made and expressed as a byproduct of making their own site useful, to produce a highly accurate relevance and accreditation algorithm. Its experience is particularly valuable when juxtaposed to that of GoTo.com, which was a search engine that specifically sold placement on the search result list to the highest bidder. It turns out, however, that one person's willingness to pay *to be seen* is not necessarily a good measure of the utility its site provided to people who are searching the web in the area that that person occupies. Google recently replaced Overture, GoTo's current name, as AOL's default search engine.⁴⁶ A casual search using both will reveal the different qualities of the two, and a search for "Barbie" will also yield interesting insights into the political morality of pricing as opposed to voting as the basis of relevance algorithms.

While Google is an automated mechanism of collecting human judgment as a by product of some other activity (publishing a web page) there are also important examples of distributed projects self-consciously devoted to peer production of relevance. Most prominent among these is the Open Directory Project.⁴⁷ The site relies on tens of thousands of volunteer editors to determine which links should be included in the directory. Acceptance as a volunteer requires application. Not all are accepted, and quality relies on a peer review process based substantially on seniority as a volunteer and engagement. The site is hosted and administered by Netscape, which pays for server space and a small number of employees to administer the site and set up the initial guidelines, but licensing is free, and presumably adds value partly to AOL's and Netscape's commercial search engine/portal, and partly through goodwill. The volunteers are not affiliated with Netscape, receive no compensation, and manage the directory out of the joy of doing so, or for other internal or external

⁴⁵ See description <http://www.google.com/technology/index.html>.

⁴⁶ David F. Gallagher, AOL Shifts Key Contract to Google, NYT (May 2, 2002), C4, col 6.

⁴⁷ <http://www.dmoz.org>.

motivations. Out of these motivations the volunteers spend time on selecting sites for inclusion in the directory (in small increments of perhaps 15 minutes per site reviewed), producing the most comprehensive, highest quality human-edited directory of the Web—competing with, and quite possibly outperforming, Yahoo in this category.

Perhaps the most elaborate mechanism for peer production of relevance and accreditation, at multiple layers, is Slashdot.⁴⁸ Billed as “News for Nerds”, Slashdot primarily consists of users commenting on initial submissions that cover a variety of technology-related topics. The submissions are typically a link to an off-site story, coupled with some initial commentary from the person who submits the piece. Users follow up the initial submission with comments that often number in the hundreds. The initial submissions themselves, and more importantly the approach to sifting through the comments of users for relevance and accreditation, provide a rich example of how this function can be performed on a distributed, peer production model.

First, it is important to understand that the function of posting a story from another site onto Slashdot, the first “utterance” in a chain of comments on Slashdot, is itself an act of relevance production. The person submitting the story is telling the community of Slashdot users “here is a story that people interested in “News for Nerds” should be interested in.” This initial submission of a link is itself filtered by “authors” (really editors) who are largely paid employees of Open Source Development Network (OSDN), a corporation that sells advertising on Slashdot and customized implementations of the Slash platform. Stories are filtered out if they have technical formatting problems or, in principle, if they are poorly written or outdated. This segment of the service, then, seems mostly traditional—paid employees of the “publisher” decide what stories are, and what are not, interesting and of sufficient quality. The only “peer production” element here is the fact that the initial trolling of the web for interesting stories is itself performed in a distributed fashion. This characterization nonetheless must be tempered, because the filter is relatively coarse, as exemplified by the FAQ response to the question, “how do you verify the accuracy of Slashdot stories?” The answer is, “We don’t. You do. If something seems outrageous, we might look for some corroboration, but as a rule, we regard this as the responsibility of the submitter and the audience. This is why it’s important to read comments. You might find something that refutes, or supports, the story in the main.”⁴⁹ In other words, Slashdot very self-consciously is organized as a means of facilitating peer production of accreditation—it is at the comments stage that the story undergoes its most important form of accreditation—peer review *ex post*.

⁴⁸ <http://www.slashdot.org>.

⁴⁹ <http://slashdot.org/faq/editorial.shtml#ed230>.

And things do get a lot more interesting as one looks at the comments. Here, what Slashdot allows is the production of commentary on a peer-based model. Users submit comments that are displayed together with the initial submission of a story. Think of the “content” produced in these comments as a cross between academic peer review of journal submissions and a peer-produced substitute for television’s “talking heads.” It is in the means of accrediting and evaluating these comments that Slashdot’s system provides a comprehensive example of peer production of relevance and accreditation.

Slashdot implements an automated system to select moderators from the pool of the users.⁵⁰ Moderators are selected according to several criteria; they must be logged in (not anonymous), they must be regular users (selects users who use the site averagely, not one time page loaders or compulsive users), they must have been using the site for a while (defeats people who try to sign up just to moderate), they must be willing, and they must have positive “karma”. Karma is a number assigned to a user that primarily reflects whether the user has posted good or bad comments (according to ratings from other moderators). If a user meets these criteria, the program assigns the user moderator status and the user gets five “influence points” to review comments. The moderator rates a comment of his choice using a drop down list with words such as “flamebait” and “informative”. A positive word increases the rating of a comment one point and a negative word decreases the rating a point. Each time a moderator rates a comment it costs the moderator one influence point, so the moderator can only rate five comments for each moderating period. The period lasts for three days and if the user does not use the influence points, they expire. The moderation setup is designed to give many users a small amount of power – thus decreasing the effect of rogue users with an axe to grind or with poor judgment. The site also implements some automated “troll filters” which prevent users from sabotaging the system. The troll filters prevent users from posting more than once every 60 seconds, prevent identical posts, and will ban a user for 24 hours if the user has been moderated down several times within a short time frame.

Slashdot provides the users with a “threshold” filter that allows each user to block lower quality comments. The scheme uses the numerical rating of the comment (ranging from –1 to 5). Comments start out at 0 for anonymous posters, 1 for registered users and 2 for registered users with good “karma”. As a result, if a user sets their filter at 1, the user will not see any comments from anonymous posters unless the comments’ ratings were increased by a moderator. A user can set their filter anywhere from –1 (viewing all of the comments), to 5 (where only the posts that have been upgraded by several moderators will show up).

⁵⁰ The description in the following few paragraphs is mostly taken from the site’s FAQ, <http://slashdot.org/faq/editorial.shtml#ed230> or from observations.

Relevance, as distinct from accreditation, is also tied into the Slashdot scheme because off topic posts should receive an “off topic” rating by the moderators and sink below the threshold level (assuming the user has the threshold set above the minimum). However, the moderation system is limited to choices that sometimes are not mutually exclusive. For instance, a moderator may have to choose between “funny” (+1) and “off topic” (-1) when a post is both funny and off topic. As a result, an irrelevant post can increase in ranking and rise above the threshold level because it is funny or informative. It is unclear, however, whether this is a limitation on relevance, or indeed mimics our own normal behavior, say in reading a newspaper or browsing a library (where we might let our eyes linger longer on a funny or informative tidbit, even after we’ve ascertained that it is not exactly relevant to what we were looking for).

The primary function of moderation is to provide accreditation. If a user sets a high threshold level, they will only see posts that are considered of high quality by the moderators. Users also receive accreditation through their karma. If their posts consistently receive high ratings, their karma will increase. At a certain karma level, their comments will start off with a rating of 2 thereby giving them a louder voice in the sense that users with a threshold of 2 will now see their posts immediately, and fewer upward moderations are needed to push their comments even higher. Conversely, a user with bad karma from consistently poorly rated comments can lose accreditation by having their posts initially start off at 0 or -1. At the -1 level, the posts may not get moderated, effectively removing the opportunity for the “bad” poster to regain karma.

In addition to the mechanized means of selecting moderators and minimizing their power to skew the aggregate judgment of the accreditation system, Slashdot implements a system of peer-review accreditation for the moderators themselves. Slashdot implements this “meta-moderation” by making any user that has an account from the first 90% of accounts created on the system eligible to moderate the moderations. Each eligible user who opts to perform meta-moderation review is provided with 10 random moderator ratings of comments. The randomness itself helps to prevent biases and control by anyone who might use the assignment process to influence the selection of moderators. The user/meta-moderator then rates the moderator’s rating as either unfair, fair, or neither. The meta-moderation process affects the karma of the original moderator, which, when lowered sufficiently by cumulative judgments of unfair ratings, will remove the moderator from the moderation system.

Together, these mechanisms allow for the distributed production of both relevance and accreditation. Because there are many moderators who can moderate any given comment, and thanks to the mechanisms that explicitly limit the power of

any one moderator to over-influence the aggregate judgment, the system evens out differences in evaluation by aggregating judgments. The system then allows individual users to determine what level of accreditation pronounced by this aggregate system fits their particular time and needs by setting their filter to be more or less inclusive. By introducing “karma,” the system also allows users to build reputation over time, and to gain greater control over the accreditation of their own work relative to the power of the critics. Users, moderators, and meta-moderators are all volunteers. The primary point to take from the Slashdot example is that the same dynamic that we saw used for peer production of initial utterances, or content, can be implemented to produce relevance and accreditation. Rather than using the full time effort of professional accreditation experts the system is designed to permit the aggregation of many small judgments, each of which entails a trivial effort for the contributor, regarding both relevance and accreditation of the materials. The software that mediates the communication among the collaborating peers embeds both the means to facilitate the participation and a variety of mechanisms designed to defend the common effort from poor judgment or defection.

3. Value-added Distribution

Finally, when we speak of information or cultural goods that exist (content has been produced) and are made usable through some relevance and accreditation mechanisms, there remains the question of “distribution.” To some extent this is a non-issue on the Net. Distribution is cheap, all one needs is a server and large pipes connecting one’s server to the world, and anyone, anywhere can get the information. I mention it here for two reasons. One, there are a variety of value-adding activities at the distribution stage—like proofreading in print publication—that need to be done at the distribution stage. Again, as long as we are talking about individual web sites, the author who placed the content on the Web will likely, for the same motivations that caused him or her to put the materials together in the first place, seek to ensure these distribution values. Still, we have very good examples of precisely these types of value being produced on a peer production model. Furthermore, as the Net is developing, the largest ISPs are trying to differentiate their services by providing certain distribution-related values. The most obvious examples are caching and mirroring—implementations by the ISP (caching) or a third party like Akamai (mirroring) that insert themselves into the distribution chain in order to make some material more easily accessible than other material.⁵¹ The question is the extent to which peer distribution can provide similar or substitute values.

⁵¹ Part of the time lag involved in downloading materials is the time it takes for the materials to traverse the network from their point of origin to the user’s computer. One approach to speeding up communications is to store copies of popular materials close to users, so that the need for large amounts of information to travel across multiple networks and servers is decreased. When Internet Service Providers do this, the function is called “caching”, which relates to temporary storage of recently viewed

The most notorious example is Napster.⁵² The point here was that the collective availability of tens of millions of hard drives of individual users provided a substantially more efficient distribution system for a much wider variety of songs than the centralized (and hence easier to control) distribution systems preferred by the recording industry. The point here is not to sing the praises of the dearly departed (as of this writing) Napster. The point is that, setting aside the issues of content ownership, efficient distribution could be offered by individuals for individuals. Instead of any one corporation putting funds into building a large server and maintaining it, end-users opened part of their hard drives to make content available to others. And while Napster required a central addressing system to connect these hard drives, Gnutella and other emerging peer-to-peer networks does not.⁵³ This is not the place to go into the debate over whether Gnutella has its own limitations, be they scalability or free riding.⁵⁴ The point is that there are both volunteers and commercial software companies involved in developing software intended to allow users to set up a peer-based distribution system that will be independent of the more commercially-controlled distribution systems, and will be from the edges of the network to its edges, rather than through a controlled middle.⁵⁵

Perhaps the most interesting, discrete and puzzling (for anyone who dislikes proofreading) instantiation of peer-based distribution function is Project Gutenberg

files. See David D. Clark & Marjorie Blumenthal, Rethinking the Design of the Internet: The end to end arguments vs. the brave new world, at 15, available <http://lawschool.stanford.edu/e2e/papers/TPRC-Clark-Blumenthal.pdf>. Akamai is a business that provides similar functionality independently of the ISP, allowing content providers to purchase the functionality independently of the decisions of an ISP. See, Akamai White Paper: Turbo Charging Dynamic Web Sites with Akamai EdgeSuite, available from http://www.akamai.com/en/resources/pdf/Turbocharging_WP.pdf. So, for example, if CNN wants to be served quickly, but AT&T Worldnet is not caching CNN, CNN can use the services of Akamai to "mirror" its site in many important local markets so that whoever accesses the materials will receive the more rapid service.

⁵² See generally <http://dir.salon.com/topics/napster/index.html> (collecting variety of stories and explanations of the rise and fall of said dearly departed)

⁵³ See Andy Oram, Gnutella and FreeNet Represent True Technological Innovation, O'Reilly Magazine May 12, 2000, available <http://www.oreillynet.com/pub/a/network/2000/05/12/magazine/gnutella.html>.

⁵⁴ See Eytan Adar and Bernardo Huberman, Free Riding on Gnutella, 5(10) First Monday (Oct. 2 2000), available http://www.firstmonday.dk/issues/issue5_10/adar/. But see Clay Shirky, In praise of free loaders, 12/01/2000, available http://www.openp2p.com/pub/a/p2p/2000/12/01/shirky_freeloading.html.

⁵⁵ Eben Moglen has argued that peer distribution has dramatically improved characteristics over proprietary distribution, because social familiarity leads people to be better at guessing their friends' and acquaintances' preference than is a centralized distributor. Giving individuals the freedom to give their friends music or any form of utterance that they believe they will like, and within a very small number of steps of social propagation the information will arrive in the hands of most everyone who would want it. Eben Moglen, The dotCommunist Manifesto: How Culture Became Property and What We're Going to Do About It", University of North Carolina, Chapel Hill, University Program in Cultural Studies, November 8, 2001, <http://www.ibiblio.org/moglen>.

and the site set up to support it, Distributed Proofreading. Project Gutenberg⁵⁶ entails hundreds of volunteers who scan in and correct books so that they are freely available in digital form. Currently, Project Gutenberg has amassed around 3,500 public domain “etexts” through the efforts of volunteers and makes the collection available to everyone for free. The vast majority of the etexts are offered as public domain materials. The etexts are offered in ASCII format, which is the lowest common denominator and makes it possible to reach the widest audience. The site itself presents the etexts in ASCII format but does not discourage volunteers from offering the etexts in markup languages. It contains a search engine that allows a reader to search for typical fields such as subject, author and title. Distributed Proofreading is a site that supports Project Gutenberg by allowing volunteers to proofread an etext by comparing it to scanned images of the original book. The site is maintained and administered by one person.

Project Gutenberg volunteers can select any book that is in the public domain to transform into an etext. The volunteer submits a copy of the title page of the book to Michael Hart—who founded the project—for copyright research. The volunteer is notified to proceed if the book passes the copyright clearance. The decision on which book to convert to etext is left up to the volunteer, subject to copyright limitations. Typically a volunteer converts a book to ASCII format using OCR (optical character recognition) and proofreads it one time in order to screen it for major errors. The volunteer then passes the ASCII file to a volunteer proofreader. This exchange is orchestrated with very little supervision. The volunteers use a listserv mailing list and a bulletin board to initiate and supervise the exchange. In addition, books are labeled with a version number indicating how many times they have been proofed. The site encourages volunteers to select a book that has a low number and proof it. The Project Gutenberg proofing process is simple and involves looking at the text itself and examining it for errors. The proofreaders (aside from the first pass) are not expected to have access to the book or scanned images, but merely review the etext for self-evident errors.

Distributed Proofreading,⁵⁷ a site unaffiliated with the Project Gutenberg, is devoted to proofing Project Gutenberg etexts more efficiently, by distributing the volunteer proofreading function in smaller and more information rich modules. In the Distributed Proofreading process, scanned pages are stored on the site and volunteers are shown a scanned page and a page of the etext simultaneously so that the volunteer can compare the etext to the original page. Because of the fine-grained modularity, proofreaders can come on the site and proof one or a few pages and submit them. By contrast, on the Project Gutenberg site the entire book is typically exchanged, or at

⁵⁶ <http://promo.net/pg/>

⁵⁷ <http://charlz.dynip.com/gutenberg/>

minimum a chapter. In this fashion, Distributed Proofreading clears the proofing of thousands of pages every month.

What is particularly interesting in these sites is that they show that even the most painstaking, some might say mundane, jobs can be produced on a distributed model. Here the motivation problem may be particularly salient, but it appears that a combination of bibliophilia and community ties suffices (both sites are much smaller and more tightly knit than, say, the Linux kernel development community). The point is that individuals can self-identify as having a passion for a particular book, or as having the time and inclination to proofread as part of a broader project they perceive to be in the public good. By connecting a very large number of people to these potential opportunities to produce, the e-text projects, just like clickworkers, or Slashdot, or Amazon, can capitalize on an enormous pool of underutilized intelligent human creativity and willingness to engage in intellectual effort.

4. Summary

What I hope these examples provide is a common set of mental pictures of what peer production looks like. In the remainder of the article I will abstract from these stories some general observations about peer production, what makes it work, and what makes it better under certain circumstances than market- or hierarchy-based production. But at this point it is important that the stories have established the plausibility of, or piqued your interest in, the claim that peer production is a phenomenon of much wider application than free software, and that it is something that actually exists. What remains is the interesting and difficult task of explaining it in terms comprehensible to those who make economic policy, so as to begin to think about the policy implications of the emergence of this strange breed in the middle of our information economy. I will by no stretch of the imagination claim to have completed this task in the following pages. But I hope to identify some basic regularities and organizing conceptions that will be useful to anyone interested in pursuing the answer. Even if you do not buy a single word of my initial efforts to theorize the phenomenon, however, seeing these disparate phenomena as instances of a general emerging phenomenon in the organization of information production should present a rich topic of study for organization theorists, anthropologists, institutional economists, and business people interested in understanding new production models in a ubiquitously networked environment.

II. Why Would Peer Production Emerge in a Networked Environment?

1. Locating the theoretical space for peer production

There are many places to locate an attempt to provide a theoretical explanation of peer production. One option would be to focus on literature concerning the organization of production that would be most sympathetic to the sustainability and productivity of peer production. This might include the literature regarding trust-based modes of organizing production⁵⁸ or literature that focuses on internal motivation and its role in knowledge production.⁵⁹ Perhaps more importantly, it makes obvious sense to focus on cultural or sociological characteristics of peer communities as a central explanation of peer production, starting with mainstream sociological and anthropological literature of gift giving and reciprocity.⁶⁰ There are applications that are rather close both online⁶¹ and offline,⁶² as well as in economic analysis of organization.⁶³ The advantage of doing so would be that these approaches have rich and detailed analytic tools with which to analyze the phenomenon of peer

⁵⁸ See Paul S. Adler, Market, Hierarchy, and Trust: The Knowledge Economy and the Future of Capitalism, *Organization Science*, 214-234 (March-April 2001).

⁵⁹ See Margit Osterloh and Bruno S. Frey, Motivation, Knowledge Transfer, and Organizational Form (1999) Institute for Empirical Research in Economics, University of Zurich, Working Paper No. 27.

⁶⁰ These studies hark back to Franz Boas's study, *The Social Organization and the Secret Societies of the Kwakiutl Indians*, in, Annual Report of the Smithsonian Institution for 1895 (1897). The late '60s produced some of the seminal works in both sociology and anthropology of gift exchange and reciprocity. See Marcel Mauss, *The Gift: the Form and Reason for Exchange in Archaic Societies*. (1990) (orig. 1967); Marshall D. Sahlins, *On the Sociology of Primitive Exchange* in *The Relevance of Models for Social Anthropology*, (Michael P. Banton, ed. 1968); Claude Levi-Strauss, *The Elementary Structures of Kinship* (1969).

⁶¹ Indeed, this is central to Raymond's discussion of open source development, See Raymond, Homesteading the Noosphere, *supra*, though it is not entirely clear that his description in fact fits the gift literature, given how distant and potentially disconnected the act of giving in open source communities is from the act of receiving.

⁶² In the offline world, the scientific community has been described as thriving, at least in part, on shared social commitments to, for example, the pursuit of truth, progress, and open collaboration. Science in particular has been the subject of sociological analysis of a productive enterprise. Classics are Bernard Barber, *Science and the Social Order* (1953); Warren O. Hagstrom, *The Scientific Community* (1965); Robert K. Merton, *The Sociology of Science* (1973). Studies of gift exchange flow from the 1925 fountainhead of Marcel Mauss, *The Gift: Forms and Functions of Exchange in Archaic Societies* (Ian Cunnison trans., 1954). Work in this vein has followed both in anthropology and sociology. For a review of this literature and its application to current debates over patenting basic research, see Arti Kaur Rai, *Regulating Scientific Research: Intellectual Property Rights and the Norms of Science*, 94 Nw. U. L. Rev. 77 (1999).

⁶³ Dan Kahan, Reciprocity. See Ernest Fehr and Klaus M. Schmidt, Theories of Fairness and Reciprocity, Evidence and Economic Applications, Institute for Empirical Research in Economics, University of Zurich, Working Paper No. 75 (February 2001); Bruno S. Frey and Stephan Meier, Pro-Social Behavior, Reciprocity, or Both? Institute for Empirical Research in Economics, University of Zurich, Working Paper No. 107 (February 2002); Ernest Fehr and Armin Falk, Psychological Foundations of Incentives, Schumpeter Lecture, Annual Conference of the European Economic Association 2001.

production. The disadvantage is that at present the study of organization and productivity, particularly in the context of discussions of law and policy, is heavily based in economics. A discussion founded largely on methodologies that are not tractable within economic theory would fail to engage that core discourse.⁶⁴ In this early study of the phenomenon of peer production it seems more important to establish its baseline plausibility as a sustainable and valuable mode of production within the most widely used relevant analytic framework than to offer a detailed explanation of its workings. Doing so should provide wider recognition of the policy implications, and create a space for more methodologically diverse inquiries.

My effort here will be directed towards offering an explanation within the framework that has largely become the mainstream economic theory of organizations, namely, the approach that followed Ronald Coase's *The Nature of the Firm* in focusing on the comparative costs of institutional alternatives as an explanation for their emergence and relative prevalence. At the most general intuitive level, we can begin by looking at Coase's explanation of the firm and Harold Demsetz's explanation of property rights.⁶⁵ Coase's basic explanation of why firms emerge—in other words, why clusters of individuals operate under the direction of an entrepreneur, a giver of commands, rather than interacting purely under the guidance of prices—is that using the price system is costly. Where the cost of achieving a given outcome in the world through the price system will be higher than the cost of using a firm to achieve the same result, firms will emerge to organize the behavior that would attain that result. Any given firm will cease to grow when the increased complexity of its organization makes its internal decision costs higher than the costs that another, smaller firm would encounter to achieve the same marginal result. Firms as a whole will cease to incorporate actions aimed to achieving an outcome into their scope once the cost of doing so exceeds the cost of achieving that result through the market. Assuming that the cost of organization increases with size, Coase posited that we have a “natural”—i.e., internal to the theory—limit on the size and number of organizations.

⁶⁴ This is not to say that there is no literature within economics that attempts to use the gift exchange literature to study economic phenomena. Examples, in addition to those cited in footnote 63 *supra*, are George A. Akerlof, *Labor Contracts as Partial Gift Exchange*, 97 *Q. J. of Econ.* 543 (1982); Rachel E. Kranton, *Reciprocal Exchange: A Self-Sustaining System*, 86 *Am. Ec. Rev.*, 830 (1996); Janet T. Landa, *The Enigma of the Kula Ring* in Trust, ethnicity, and identity: beyond the new institutional economics of ethnic trading networks, contract law, and gift-exchange 141-172 (1994); Ernst Fehr, Erich Kirchler, Andreas Weichbold, & Simon Gächter, *When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets*, 16 *J. Labor Econ.* 324 (1998). My point is narrower, that is, that the baseline response of most economists and lawyers trained to look at questions through an economic prism is disbelief, and my purpose in the analytic segment of this article is to respond to that widely held disbelief.

⁶⁵ Harold Demsetz, *Toward a Theory of Property Rights*, 57 *Am. Econ. Rev.* 347-357 (1967).

Demsetz's basic explanation of why property emerges with regard to resources that previously were managed without property rights—as commons—can be resolved to a very similar rationale. As long as the cost of implementing and enforcing property rights in a given resource is higher than the value of the total increase in the efficiency of utilization of the resource gained by the introduction of a property regime where none existed before, the resource will operate without property rights. Once the value of the resource increases due to an exogenous circumstance—a technological development or an encounter with another civilization—so that intensification of its utilization through property-based appropriation is worth the cost of implementing property rights, property rights emerge. More generally, property in a resource emerges if the social cost of having no property in that resource exceeds the social cost of implementing a property system in it. This restatement can include within it common property regimes, managed commons, and other non-property approaches to managing sustainable commons.⁶⁶

Table 1 tabulates the interaction between Coase's theory of the firm and Demsetz's theory of property.

	Property more valuable than implementation costs ⁶⁷	Cost of implementing property higher than opportunity cost of property*
Market exchange of x more efficient than organizing x	Pure market	Pure commons? ("market" with payment by time and effort)
Organizing x more efficient than market exchange of x	Firms	Common property regimes (if organizing valuable)

Table 1: Organizational forms as a function of relative social cost of property vs. no-property and firm-based management vs. market

* Both markets and firms generally rely on property rights. The institutions described in this column, where property is assumed not to emerge, therefore reflect functional equivalents for decentralized (market) and coordinated organization (firms) in the absence of property. These are described below on pages 35-36.

Before going in to *why* peer production may be less costly than property/market-based production or organizational production, it is important to recognize that if we posit the existence of such a third option it is relatively easy to adapt the transactions cost theory of the firm and the comparative institutional cost theory of property to include it. We would say that when the cost of organizing an

⁶⁶ For discussions of commons see, Carol Rose, *The Comedy of the Commons: Custom, Commerce, and Inherently Public Property*, 53 U. Chi. L. Rev. 711 (1986); Elinor Ostrom, *Governing the Commons* (1992). A brief discussion of these concepts as applied to peer production follows below, pages 35-36.

⁶⁷ "Valuable" as compared to the option, and opportunity costs, of not having property rights in place.

activity on a peered basis is lower than the cost of using the market, and the cost of peering is lower than the cost of hierarchical organization, then peer production will emerge.⁶⁸

We could tabulate as follows:

	Property more valuable than implementation costs	Cost of implementing property higher than opportunity cost of property
Market exchange of x more efficient than organizing/ peering x	Pure market (farmers markets)	Pure commons (ideas & facts; highways?)
Organizing x more efficient than market exchange or peering of x	market with firms	Common property regimes (Swiss pastures)
Peering more efficient than market exchange or organization of x	Proprietary "open source" efforts (Xerox's Eureka)	Peer production processes ⁶⁹ (free software; academic science; NASA clickworkers)

Table 2: Organizational forms as a function of relative social cost of property vs no-property and firm-based management vs. market vs. peering

Understanding that in principle the same framework that explains the emergence of property and firms could explain the emergence of peer production focuses our effort on trying to understand why it is that peering could, under certain circumstances, be a more cost effective institutional form than either markets or hierarchical organizations. Because the emergence of peer production seems to be tied to the emergence of a pervasively networked information economy, my explanation seeks to be (1) in some sense sensitive to changes in the nature of the human and material resources used in information production relative to other productive enterprises, and (2) affected by the cost and efficiency of communication among human participants in the productive enterprise.

2. Peer production of information in a pervasively networked environment

Peer production is emerging as an important mode of information production because of four attributes of the pervasively networked information economy. First,

⁶⁸ In the context of land, Ellickson extends Demsetz's analysis in precisely this fashion, suggesting that there may be a variety of reasons supporting group ownership of larger tracts, including the definition of efficient boundaries (efficient for the resource and its use), coping with significant shocks to the resource pool, risk spreading, "the viability of group ownership might be enhanced by the advent of inexpensive video cameras or other technologies for monitoring behavior within a group setting." Robert Ellickson, *Property in Land*, 102 Yale L.J. 1315, 1330 (1993).

⁶⁹ "Cost" here would include the negative effects of intellectual property on dissemination and downstream productive use.

the object of production—information—is quirky as an object of economic analysis, in that (a) it is purely non-rival⁷⁰ and (b) its primary non-human input is the same public good as its output—information.⁷¹ Second, the physical capital costs of information production have declined dramatically with the introduction of cheap-processor-based computer networks. Third, the primary human input—creative talent—is highly variable, more so than traditional labor, and certainly more so than many material resources usually central to production. Moreover, the individuals who are the “input” possess better information than anyone else about the variability and suitability of their talents and level of motivation and focus at a given moment to given production tasks. Fourth and finally, communication and information exchange across space and time are much cheaper and more efficient than ever before, which permits the coordination of widely distributed potential sources of creative effort and the aggregation of actual distributed effort into usable end products.

The first attribute—the public goods nature of information—affects the cost of one major input into production—existing information. It means that the social cost of using existing information as input into new information production is zero.⁷² This has two effects on the relative cost of peer production of information.⁷³ First, it lowers the expected social cost of peer production of *information*, as compared to normal economic goods, because in principle it means that a central input—pre-existing information—could be available to human productive agents without limit, if the provisioning problem can be solved without introducing appropriation. Second, it underlies a pervasive social cost of market and hierarchy in this field of production, because of the losses in both static and dynamic efficiency entailed by the

⁷⁰ A good is nonrival to the extent that its consumption by one person does not diminish its availability for use by any other person. See, e.g., Paul Romer, Endogenous Technological Change, 98(5) *Journal of Political Economy*, S73-S74 (1990). It has been commonplace for a long time to treat information as a perfectly non-rival good, see *id.*; Kenneth J. Arrow, Economic Welfare and the Allocation of Resources for Invention, in *The Rate and Direction of Inventive Activity: Economic and Social Factors* 609, 616-17 (National Bureau of Economic Research, 1962) (hereinafter “Arrow 1962”).

⁷¹ While the input characteristic of information has been appreciated at least since Arrow 1962, the extensive exploration of the implications of this characteristic largely begins with See Suzanne Scotchmer, Standing on the Shoulders of Giants: Cumulative Research and the Patent Law, 5 *J. of Economic Perspectives* 29-41 (1991).

⁷² Saying that the social cost of its use by one person is zero is simply another way of saying that the good is nonrival, that is, that its use by one person does not prevent its use by any other person.

⁷³ The public goods problem of information production is usually described as comprised of two distinct characteristics, its nonrivalry and its nonexcludability. See *Romer, supra* note __, at S73-S74. A good is excludable to the extent that its producer can exclude others from its use, unless they pay. If a good is not excludable, it too presents a problem for market provisioning, but not because it is inefficient to price it positively, but because it is difficult to do so, and hence firms will provide too little of it. Non excludability of information is less important to the analysis here, because it does not relate to the characteristics of information that are important to making peer-production both feasible and efficient: that is, that its most efficient price is zero, and that it can be used by any number of people without diminishing its availability for others.

implementation of property rights in a nonrival public good usually thought necessary to sustain market and hierarchy-based production of information.⁷⁴

The second attribute—the decline in the capital cost of information production—similarly lowers the cost of another major capital cost of information production. The age of mechanical reproduction that enabled both fixation and distribution of information and culture as goods was defined by the high cost of the physical capital necessary for mechanical reproduction. The large circulation automated printing presses, vinyl record and later CD manufacturing facilities, and the movie studios and their celluloid-based systems formed the basis for the industrial model typical of information and cultural production in the 20th century. The declining cost of computer processors, however, coupled with the digitization of the fixation and transmission of all forms of information and culture, have made the physical capital necessary to embody and disseminate such goods cheaper by orders of magnitude than in the past.

Together, these first two attributes make information production a potentially sustainable low-cost, low returns endeavor for many individuals relying on indirect appropriation.⁷⁵ It is important to note, however, that the public goods attribute limits the applicability of my observations about peer production, so that I make no claim about the applicability of these observations to traditional economic goods.

The third characteristic—the centrality of human capital to information production and its variability—is, as I will explain below, the primary source of efficiency gains from moving from markets or hierarchical organization to peering. Peer production better produces information about available human capital, and increases the size of the sets of agents and resources capable of being combined in projects—where there are increasing returns to scale, in terms of allocation efficiency, for these sets.

⁷⁴ See Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, in *The Rate and Direction of Inventive Activity: Economic and Social Factors* 609, 617 (National Bureau of Economic Research, 1962) (“precisely to the extent that [property rights in information] are successful, there is an underutilization of the information”).

⁷⁵ “Indirect appropriation” is appropriation of the value of one’s effort by means other than reliance on the excludability of the product of the effort. So, someone who is paid as a teacher, but gets the position in reliance on his scholarship, is indirectly appropriating the benefit of his scholarship. An IBM engineer who gains human capital by working on GNU/Linux from home in the evening is indirectly appropriating the benefits of her efforts in participating in the production of GNU/Linux. The term is intended to separate out appropriation that is sensitive to excludability of information, which is what intellectual property is aimed to provide, and which I call direct appropriation, and appropriation that is independent of intellectual property because it is not based on exclusion from the information itself, which I call indirect appropriation. See Benkler, *Intellectual Property and the Organization of Information Production*, *supra* note ___, at 87. As a general matter, the more a sector of information production can be sustained through indirect appropriation the less it needs intellectual property.

The fourth attribute—the dramatic decline in communications costs—radically reduces the cost of peering relative to what was possible in the material world. This allows substantially cheaper movement of information inputs to human beings and of human talent to resources, and movement of modular contributions to projects, so that widely dispersed contributions can be integrated into finished information goods. It also allows communication among participants in peer production enterprises about who is doing what, and what needs to be done.

3. Markets, hierarchies, and peer production as information processing systems

Usually the question of why anyone would contribute to a peer production enterprise without directly appropriating the benefits is foremost in people's minds when I describe the phenomenon. For the sake of completeness of the organization theory argument, however, suspend disbelief for one more section (or if you cannot do so, read Part III first, and then come back here.) Assume for the next ten pages that I have come up with a reasonably plausible story as to why people participate. Doing so will allow us to consider the claim that if they did, their efforts would be more productive than if they organized in a firm or interacted purely through a price system.

Peer production has a relative advantage over firm or market-based production along two dimensions, both a function of the highly variable nature of human capital. The first emerges when one treats all approaches to organizing production as mechanisms by which individual agents reduce uncertainty as to the likely value of various courses of productive action.⁷⁶ Differences among these modes in terms of their information processing characteristics could then account for differences in their relative value as mechanisms for organizing production. The second dimension is that a particular strategy that firms, and to a lesser extent markets, use to reduce uncertainty—securing access to limited sets of agents and resources through contract and property—entails a systematic loss of productivity

⁷⁶ That is, individuals who are presented with alternative things that they might do, say, standing in a particular spot and turning a lever all day, or writing an economic analysis of friendship, do not always know which of these courses of action is more valuable, or which would allow them to put dinner on the table. One can think of markets and firms as means by which individuals solve this lack of knowledge, because they use different signals that the market or the firm generates as information about which action will better fulfill their purposes—be they glory or subsistence. What follows, then, is in some measure a sketchy application of the Herbert Simon's statement, "It is only because individual human beings are limited in knowledge, foresight, skill, and time that organizations are useful instruments for the achievement of human purpose." Simon, *Models Of Man* 199 (1957). This individual-centric view of organization, then, diverges from the traditional framework for the theory of the firm, in that the firm here is explained in terms solely related to the question of why agents use this form of organization to order their individual productive behavior. I do not differentiate between entrepreneurs, managers, and employees, but rather treat all of them as agents who have a set of possible open courses of action.

relative to peer production. This is so because there are increasing returns to scale to the size of the sets of agents and resources capable of being applied to sets of projects in terms of allocation efficiency, and peer production relies on unbounded access of agents to resources and projects.

a. Information gains

We could reduce the decisions that must be made by productive human beings as follows. Imagine a human agent, A , who is deciding whether and how to act, where act a is part of the set $\{a_1, a_2, a_3, \dots a_n\}$. Act a is a combination of two elements: the effort to be exercised, where effort e is part of the set $\{e_1, e_2, e_3, \dots e_n\}$, representing different levels and focuses of effort possible for A , and the resources as to which the effort is exerted, where resource r is part of the set $\{r_1, r_2, r_3, \dots r_n\}$ available for A to use. Both e and r are sensitive to the costs of collecting information. The components of either set are a function of the set of opportunities to exert effort and the set of resources available to work with that the agent perceives to be open to him (which are a subset of the set of resources and efforts that are open to them in principle, assuming perfect information). Both sets increase as information collection costs decline, because agents see more of the universe of opportunities actually available to them.

Imagine that A is a rational actor,⁷⁷ where the private value V_A to A of doing a is the expected value of outcome O , which is the value to A of O obtaining, discounted by the probability that O will obtain if A does a . This means that the value to A of doing a increases as the probability that doing a will result in O obtaining increases. That probability depends on the effort A will exert, the resources available to A , A 's talent t , where talent describes relative capabilities, associations, and idiosyncratic insights and educational mixes of an individual that make that person more or less productive with a given set of resources for a given project, the presence of complementary actions by other agents, a_x , and the absence of undermining actions by other agents, a_y . $V_A = qO$, where q describes the probability that $(e, r)t, a_x/a_y = O$.

A will $a_n, (e_n, r_n)$, if A believes the value V_{An} to be higher than either inaction or an alternative action. This requires that V_{An} be positive relative to the value of inaction and higher than the value V_{Am} (the value of any other O_m similarly discounted

⁷⁷ This quasi-formal statement assumes a rational actor in the most traditional sense, and can be formalized within a framework that strictly orders the value of outcomes, and with agents who know the values of their preferences and their outcomes, and can calculate the probabilities of outcomes vis-à-vis actions etc. For the limited purposes of comparing the information processing characteristics of firms, markets, and peer production, there is no necessity for these rather strong characteristics. It is enough to have individuals that, in Simon's terms, are satisficers. Simon, *Models of Man*, *supra*. All they need is that the uncertainty as to the relative value of a given action for them be reduced to a level that satisfies their own sense of the level of certainty they require to justify action, without needing fully to calculate the various outcomes.

by the probability that any other a_m , combining any (e_m, r_m) , will lead to O_m obtaining). It is important to underscore that the probability of O obtaining is in some measure dependent on the actions, both complementary and undermining, of other agents. Assuming that the agent knows his or her own valuation of the outcome, has some experience-based evaluation of t , and controls his or her own effort, e , uncertainty resides primarily with regard to the divergence of the private valuation of the outcome from its valuation by others, the availability of r , and the interdependence of the agent's action on the action or inaction of others whose actions A_n cannot control. Reducing these uncertainties is a central function of markets and firms. Reducing the latter two in particular is a central function of property and contract, which can secure complementary material and human resources and be used to increase the probability that complementary actions will be taken and decrease the probability that undermining actions will occur or be efficacious in the relevant resource set.

Markets and firm-based hierarchies are information processes in the sense that they are means of reducing the uncertainty that agents face in evaluating different courses of action to a level acceptable to the agent as a level of uncertainty warranting action. Markets price different levels of effort and resources so as to signal the relative values of actions in a way that allows individuals to compare actions and calculate the likely actions of other individuals—whose actions affect the value of the agent's action—faced with similar pricing of alternative courses of action. Firms reduce uncertainty by specifying to some individuals what actions to take, reducing uncertainty of interdependent action by controlling enough resources and people by contract and property to reduce the uncertainty of the outcomes of specified actions to a level acceptable to the managers.

To compare modes of organizing production as information processing systems one might use the term *information opportunity cost*. I use the term “information” here in the technical sense of a reduction in uncertainty, where “perfect information” is the condition where uncertainty regarding an action could not in principle be further reduced. Perfect information is impossible to come by, and different organizational modes have different strategies for overcoming this persistent uncertainty. These different strategies differ from each other in the amount or kind of information they lose in the process of resolving the uncertainty that rational agents face in deciding what course of action they should follow under given circumstances. The divergence of each mode from the hypothetical condition of perfect information is that mode's information opportunity cost.

Markets reduce uncertainty regarding allocation decisions by producing a signal that is clear, univocal (i.e., comparable across different uses) as to which use of the relevant factors would be most efficient. To do so, they require a codification of

the attributes of different levels of effort, different kinds of resources, and different attributes of outcomes, so that these can all be specified as contract terms to which a price is affixed. An example of this was the introduction of codified standards for commodities as an indispensable element of the emergence of commodities markets.⁷⁸

Since we are concerned with individual agents' decisions, and levels and focuses of effort are a major component of individual action, it is intuitive that specification and pricing of all aspects of individual effort—talent, motivation, workload, and focus as they change in small increments over the span of an individual's full day, let alone months—is impossible.⁷⁹ What we get instead is codification of effort types—a garbage collector, a law professor—that are priced more-or-less finely. But one need only look at the relative homogeneity of law firm starting salaries as compared to the high variability of individual ability and motivation levels of graduating law students to realize that pricing of individual effort can be quite crude. Similarly, these attributes are also difficult to monitor and verify over time, though perhaps not quite as difficult as predicting them *ex ante*, so that pricing continues to be a function of relatively crude information about the actual variability among people. More importantly, as aspects of performance that are harder fully to specify in advance or monitor—like creativity over time given the occurrence of new opportunities to be creative—become more important, market mechanisms become more lossy.

Markets are particularly good at resolving the uncertainties with regard to the difference in valuation of the outcome among different agents, so that an agent acting on a market price will have a relatively certain evaluation of the external valuation of the outcome. Of course, this valuation may be flawed because of externalities not reflected in the price, but (for better or worse, depending on the magnitude and shape of externalities) markets plainly reduce uncertainty about the value of an action as perceived by others. Markets reduce the uncertainty about the availability of resources similarly, by allowing an agent to compare the value of an outcome to the price of necessary resources. Finally, markets reduce uncertainties with regard to the actions of other agents in two ways. First, agents can evaluate the risk that others will act in a way that is detrimental, or fail to act in a way that is complementary to, the agent's action, given the relative pricing of the courses of complementary or detrimental action. This risk assessment can then be built into the perceived value of a possible action. Second, agents can maintain property rights in resources and

⁷⁸ See Alfred Chandler, *The Visible Hand* 211 (1977) (describing how commodity attributes became codified, and local variability squelched, as a part of the transition to commodity markets).

⁷⁹ In the context of the market for labor, this has sometimes been called the multi-task problem (the problem of inability contractually to specify completely all the tasks required and attributes of an employee who will likely need to perform multiple tasks.) See Bengt Holmstrom, *The Firm as a Sub-Economy* (1999) 15 *J.L., Econ. & Org.* 74 (1999).

projects,⁸⁰ and contract with other individuals for specific complementary services, so as to prevent opportunities for negatively correlated action and to provide relatively secure access to resources for complementary action.

Firms or hierarchical organizations resolve uncertainty by instituting an algorithm or process by which information about which actions to follow is ordered, so that some pieces of information count as a sufficient reduction in uncertainty about the correct course of action to lead agents who receive them to act. The mythical entrepreneur (or the historical manager) becomes the sole source of information that is relevant to reducing the uncertainty of the workers in a purely managed firm. In the ideal-type firm, the question of incentives—reducing uncertainty as to which of a set of actions will increase an agent’s welfare—is reduced not by reference directly to market signals, but by fixing a salary for a stated behavior (following a manager’s orders) and shifting some of the risk of that course of action from employees to employers. Production processes (if I stand here and twist this lever all day, cars will emerge from the other side and I will get a paycheck) are codified as instruction sets. Agents reduce their uncertainty as to why to act and what to do by reducing the universe of information they deem relevant to their decision. Information that arrives in a particular channel with a particular level of authorization counts as signal, and all the rest as noise. It remains to the entrepreneur (in the pure model of the firm)⁸¹ to be the interface between the firm and the market, and to translate one set of uncertainty reducing signals—prices—to another set of signals with similar effect—organizational commands.

By controlling a set of resources through property and commanding a set of agents through the employment relationship, the firm reduces the elements of uncertainty related to the interdependence of the actions of agents. But by doing so it

⁸⁰ Maintaining rights in what I call “projects” is, on Kitch’s now-classic reading, the primary function of the patent system. See Edmund Kitch, *The Nature and Function of the Patent System*, 20 J.L. & Econ. 265 (1977). Even if one is critical of Kitch’s almost-exclusive focus on this characteristic as the reason for the patent system, recognizing that in some measure patents provide control over projects is all that is necessary here. The derivative use right in copyright plays a similar function to a more limited extent.

⁸¹ A number of readers have complained that this picture of the firm is too thin to be a realistic complete description. Firms use all sorts of market-based mechanisms like incentive compensation and internal “arms length” bargaining among units, and mix market-based and hierarchical control mechanisms to organize production. The point of my description, however, is not to present a true sociological description of production in a firm. The transaction costs theory of the firm identifies two dimensions to the process of allocating resources—pricing and managerial commands—making it possible to map different organizations according to whether and how they mix these ideal type modes of coordinating the use of resources in production. I present these two ideal types in their ideal form here, so as to clarify what is different about peer production within this theory. In effect, within this theory, peer production would emerge as a third dimension, to create a three-dimensional space within which an organizational model can be described. In this model, employee-of-the-month programs and employee feedback sessions become simple precursors to hybrids between firms and peer production processes, the most obvious example of which is presented by Xerox’s Eureka.

creates a boundary around the set of available agents and the set of available resources, and limits the information available about what other agents could have done with these same resources, or what else these agents could have done with these or other resources. This boundary then limits the efficacy of information collection mechanisms—like incentive-based contracts—that firms use to overcome the difficulty of collecting information to which their employees have special access.⁸² These mechanisms mean that the employees and resources within the boundary are likely to be better allocated than in firms with no similar mechanism. But firms still lose information about what human agents outside the firm could have done with these resources, or what agents within the firm could have done with resources outside the firm.

The point to see is that like the price system, hierarchical organization is a lossy medium. All the information that could have been relevant to the decision regarding each factor of production, but that was not introduced in a form or at a location that entitled it to “count” towards an agent’s decision, given the algorithm used by the organizational structure, is lost. Much of the knowledge management movement in business schools and punditry since the mid-1990s was concerned with mitigating the lossiness of managerial hierarchy as an information processing mechanism.⁸³ Mitigating this lossiness is the primary job of CIOs.⁸⁴

An example where peer production—proprietary, not commons-based—was used precisely to solve the lossiness of hierarchical organization is Xerox’s Eureka system for organizing the flow of questions from and answers to field technicians about failures of photocopiers.⁸⁵ The firm created a decentralized communications system for technicians to post questions, a peer-reviewed system for other technicians to answer these questions, and a database library of past questions and answers available to technicians who confront new problems. The original approach towards technical failures of machines was by reference to manuals that came with the machines. The machine was conceptualized as completely engineered by the engineers, with all the possible failures specified in the manual. Technicians were then conceived of as instruction followers, who came to machines that were broken, diagnosed a problem by locating it in the manual, and then executing a series of

⁸² Another problem with incentives based contracts is that they may lead to undermining voluntary cooperation, a phenomenon related to the relationship between the presence of money and social psychological rewards discussed *infra*, notes 99-103. See Ernst Fehr and Simon Gächter, *Do incentive Contracts undermine cooperation*, Working Paper No. 34, IERE (Zurich, April 2002).

⁸³ For a range of definitions of knowledge management, and a taste of the analytic styles, see <http://www.brint.com/km/kmdefs.htm>.

⁸⁴ CIOs are “Chief Information Officers” the position created to reduce information loss within organizations.

⁸⁵ See Daniel G. Bobrow, Robert Cheslow, and Jack Whalen, *Community Knowledge Sharing in Practice*, available <http://jonescenter.wharton.upenn.edu/VirtualCommunities/whalen.pdf>.

corrective steps prescribed by the manual. The Eureka project changed the conception of the knowledge content of the machine and the organizational role of technicians from instruction followers to knowledge producers. In this system, a technician who came across a problem not clearly resolved in the manual posted a question electronically to a proprietary communications system accessible by all technicians. Any other technician in the system who had come across a similar problem could post a fix, which would be reviewed by experienced technicians, who would opine on its advisability. The technician who uses the fix can then report on whether it worked. The whole transaction is stored in a database of solutions. The technicians were not compensated for answering queries, but instituted instead a system of authorship and honor-based payoffs. Eureka flipped the traditional hierarchical conception of knowledge in a machine as codified by engineers and implemented by instruction-following technicians. The knowledge content of the machine was now understood to be something that is incomplete when it leaves the design board, and is completed over the life of the machines by technicians who share questions and solutions on a peer-review, volunteer model. The Eureka project also suggests one important interpretation of peer production in relationship to markets and hierarchies. We generally understand the existence of markets and hierarchies as two ideal models of organization, and observe various mixes of the two types in actual organizational practices. Eureka suggests that peer production can be a third ideal type organizational model, which, like the other two, can be combined in various measures with the other two, forming a three dimensional map of practicable organizational strategies rather than the two dimensional map recognized traditionally.

Recognizing the lossiness of markets and managerial hierarchies suggests a working hypothesis about why peer production has succeeded in gaining ground, based on the possibility that peer production may have lower information opportunity costs than markets or hierarchies. In particular, I suggest that the primary source of gains—which may be called *information gains*—that peer production offers is its capacity to collect and process information about human capital. The hypothesis is that rich information exchange among large sets of agents free to communicate and use existing information resources cheaply will create sufficiently substantial information gains of this sort that, together with the allocation gains that I will discuss in the following section, overcome the added information exchange costs necessary to overcome the absence of pricing and managerial direction, and the added coordination costs created by the lack of property and contract as institutional bases for structuring coordination.

Where the physical capital costs of information production are low, and where existing information resources are freely or cheaply available,⁸⁶ the low cost of

⁸⁶ By limiting the hypothesis to information production under conditions of cheap and widely available physical capital (computers) and relatively free availability of information inputs, we can largely ignore

communication among very large sets of agents allows agents to collect information about available resources; possible courses of action; possible outcomes and their valuation; and the likely behavior of other agents, through extensive communication and feedback instead of by using information compression mechanisms like prices or managerial instructions. If communications include a sufficiently large number of agents operating in the same resource and opportunity sets, this mode of communication can provide to each agent rich information about what needs to or can be done, who is doing what, and how other people value any given outcome. One sees this phenomenon in the centrality of effective communications platforms to the design of peer production processes—be they the lists that lie at the heart of every free software development project⁸⁷ or the sophisticated collaboration platforms that underlie projects like Slashdot or Kuro5hin. The value of these systems is precisely in enabling agents to use extensive information exchange and feedback to provide the same desiderata that prices and managerial commands provide in their respective models. The platform design and maintenance, and more importantly the human attention required to take in and use this information, are the equivalent for peer production of organization/decision costs in firms and of transaction costs in markets. This rich information exchange may or may not be efficient, depending on the magnitude of the cost and the relative information gains generated by the richer information available to agents through this system. These, in turn, will partly be a function of the quality of the design of the collaboration platform in terms of efficiency of communication and information processing utilities.

Reducing uncertainty as to the availability of opportunities for action by any given agent and about complementary actions by other agents becomes the salient potential source of information gain for peer production projects, while the capacity of a project to reduce the likely prevalence or efficacy of undermining actions becomes a major limiting factor. This latter effect, most obviously typified by the information-rich process of peer review, will occupy a substantial portion of Part III, where I will discuss in some detail the threats to effective peer production and the mechanisms available to this mode of production to defend itself from incompetence and defection. Here I will focus on the information gains generated by peer production in terms of both opportunities for creative utilization of the existing resources in ways not

uncertainty as to availability of material resources, because the domain of application of the hypothesis relates to conditions where resources other than human creativity are not scarce, so that uncertainty as to their availability is minimal.

⁸⁷ At the heart of the distributed production system that is typified by open source software development is the notion of making the program available in a publicly accessible space for people to comment on and upgrade. See Raymond, Cathedral and Bazaar, *supra*, at 7-9. These communication lists have also offered a valuable location for observers of the phenomenon to gain insight. See, e.g., Karim Lakhani and Eric von Hippel, How Open Source Software Works: Free User to User Assistance; (Working Paper 2000), available <http://opensource.mit.edu/papers/lakhanivonhippelusersupport.pdf> (describing support lists for Apache).

previously done,⁸⁸ and the opportunities agents have to use their own talents, availability, focus, and motivation to perform a productive act.

Central to my hypothesis about the information gains of peer production is the claim that human intellectual effort is highly variable and individuated. People have different innate capabilities, personal, social, and educational histories, emotional frameworks, and ongoing lived experiences, which make for immensely diverse associations with, idiosyncratic insights into, and divergent utilization of existing information and cultural inputs at different times and in different contexts. Human creativity is therefore very difficult to standardize and specify in the contracts necessary for either market-cleared or hierarchically organized production. As human intellectual effort increases in importance as an input into a given production process, an organization model that does not require contractual specification of the effort required to participate in a collective effort and allows individuals to self-identify for tasks will be better at gathering and utilizing information about who should be doing what than a system that does require such specification. Intra-firm hybrids, like incentive compensation, may be able to improve on firm-only or market-only approaches, but it is unclear how well they can overcome the core difficulty, that is, that both approaches require significant specification of the object of organization and pricing—in this case, human intellectual input.

The point here is qualitative. It is not only, or even primarily that more people can participate in production. It is that the widely distributed model of information production will better identify who is *the best person* to produce a specific component of a project, *all abilities and availability to work on the specific module within a specific time frame* considered. With enough uncertainty as to the value of various productive activities and enough variability in the quality of both information inputs and human creative talent vis-à-vis any set of production opportunities, coordination and continuous communications among the pool of potential producers and consumers can generate better information about the most valuable productive actions, and the best human inputs available to engage in these most valuable actions at a given time. While markets and firm incentive schemes are aimed at producing precisely this form of self-identification, the rigidities associated with collecting and comprehending bids from individuals through these systems (*i.e.* transaction costs) limit the efficacy of self-identification, relative to a system where, once an individual self-identifies for a task, he or she can then undertake it without permission, contract, or instruction from another.

⁸⁸ This is a point Bessen makes about complex software, see Bessen, *supra*, as well as a characteristic of the motivation Raymond describes as having an itch to scratch. Raymond, Cathedral and Bazaar *supra*, at 4 (“Every good work of software starts by scratching a developer’s personal itch”).

Now, self-identification is not always perfect, and some mechanisms used by firms and markets to codify effort levels and abilities—like formal credentials—are the result of experience with substantial errors or misstatements by individuals of their capacities. To succeed, therefore, peer production systems must also incorporate mechanisms for smoothing out incorrect self-assessments—as peer review does in traditional academic research or in the major sites like Slashdot or Kuro5hin, or as do redundancy and statistical averaging in the case of NASA clickworkers. In information terms, these mechanisms reduce the uncertainty associated with the likely presence of undermining actions by other agents. The prevalence of misperceptions that agents have about their own ability and the cost of eliminating such errors will be part of the transaction costs associated with this form of organization that are parallel to quality control problems faced by firms and markets. This problem is less important where the advantage of peer production is in acquiring fine-grained information about motivation and availability of individuals who have otherwise widely available capabilities—like the ability to evaluate the quality of someone else's comment on Slashdot. It is likely more important where what is necessary is a particular skill set that may not be widespread—like the programming skills necessary to fix a bug in a program.

b. Allocation gains

In addition to its potential information gains, peer production has potential allocation gains enabled by one of the differences in how peer production, firms, and markets, reduce uncertainty. This gain is cumulative to the general information-processing characteristics of peer production and is a function of the size of the pools of individuals, resources, and projects engaged in information production. Like the information gains, it is based on the high variability of human capital, which suggests that there are increasing returns to the scale of the pool of individuals, resources, and projects to which they can be applied.

Market- and firm-based production rely on property and contract to secure access to bounded sets of agents and resources in the pursuit of specified projects. (As illustrated in Figure 1.) The permeability of the boundaries of these sets is limited by the costs of making decisions in a firm about adding or subtracting a marginal resource, agent, or product, and the transaction costs of doing any of these things through the market. Peer production relies on making an unbounded set of resources available to an unbounded set of agents, who can apply themselves towards an unbounded set of projects and outcomes. The variability in talent and other idiosyncratic characteristics of individuals suggests that any given resource will be more or less productively used by any given individual, and that the overall productivity of a set of agents and a set of resources will increase when the size of the sets increases towards completely unbounded availability of all agents to all resources

for all projects. The point is that even if in principle we have information as to who was the best person for a job given any particular set of resources and projects (in other words, if the information gains are assumed away), the transaction or organizational costs involved in bringing that agent to bear on the project may be too great relative to the efficiency gain over use of the resource by the next-best available agent who is within the boundary.

Assume that the productivity (P) of a set of agents/resources is a function of the Agents (A) available to invest effort (e) on resources (r). P_{A_i} (Productivity of agent A_i)⁸⁹ is a function of set of resources A_i can work on, r_i , the level of effort A_i will invest, e_i , and A_i 's talent, (t_i). P_A increases as a function of e , r , and the actions of other complementary agents x (undermined, however, by agents y , representing negative contributions), at a magnitude that is a function of t . t is a personal characteristic of individuals that is independent of the set of resources open for A to work on, but will make a particular A more or less likely to be productive with a given set of resources in collaboration with other agents for the achievement of a set of outcomes. A 's access to any given set of resources and potential collaborators therefore represents a probability, that is a function of t , that A will be productive with that resource for a given project.⁹⁰ P_A is a function of $(e, r, a_x)t/a_y$.



It is the existence of t that generates the increasing returns to the scale of the set of resources and to the set of agents to which it is available, because the larger the number of agents with access to a larger number of resources the higher the probability that the agents will include someone, A_n , with relatively high value of t_n , where t_n describes a particularly high capacity for productive use of a given r_n at a given e_n as compared to agents. If A_1 who works for firm F_1 has a higher t value as regards using r_2 than A_2 who works for F_2 , but r_2 is owned by F_2 , r_2 will nonetheless be used by A_2 so long as the value of A_2 working on r_2 has a value of no less than the value of A_1 working on r_2 minus the transaction costs involved in identifying the relative advantage of A_1 and assigning A_1 to work on r_2 . This potential efficiency loss would be eliminated if A_1 were in the set of agents who had transaction-free access to work on the resource set that includes r_2 .

So, if firms F_1 and F_2 each has a set of agents and resources, $\{A_{F1}, r_{F1}\}$ $\{A_{F2}, r_{F2}\}$, then $P_{F1} + P_{F2} < P_{F1+2}$.

⁸⁹ While V_A discussed in the previous section related to the private value of an action to an agent, $P_{A1...n}$ is intended to represent the potential social value of the efforts of any one or more agents A as part of a potential collaborative effort.

⁹⁰ In seeking to identify the private value of an outcome to an agent above, I described the successful completion of a project as an outcome O , and its value as discounted by the probability q , of O obtaining should A do a . P_A is the social equivalent of qO to the individual, representing a judgment of whether an individual will be productive.

Figures 1 and 2 illustrate the point. Figure 1 assumes that there are two firms, each having contracts with a set of agents and property in a set of resources. Assume that as among $\{A_1, \dots, A_5\}$ the best agent for using the combination r_1, r_4 is A_2 . Assume also that as among the agents $\{A_1, \dots, A_9\}$, A_8 is the best, in the sense that if A_8 were to use these resources, the social value of the product would be greater by some measure m than when, A_2 , the best agent within Firm A, uses them.





Not only is it unlikely that the two firms will have the information that A_8 is best for the job, as I suggested in the discussion of information gains. Even if they do know, as long as transaction costs associated with transferring the creativity of A_8 to Firm A or the property in r_1 and r_4 to Firm B are greater than m , creativity will be misapplied. When the firms merge, or when the agents and resources are in a common peer production enterprise space, the best person can self-identify to use the resources. Think of this as someone musing about fairy tales and coming up with a biting satire, which she is then capable of implementing, whereas the employee of the initial owner of the rights to the fairy tale might only produce a depressingly earnest new version.

This initial statement is a simplification and understatement of the potential value of the function by which the sizes of the sets of agents and resources increase productivity. There are two additional components: the range of projects that might be pursued with different talent applied to a given set of resources, and the potential for valuable collaboration. First, a more diverse set of talents looking at a set of resources may reveal available *projects* that would not be apparent when one only considers the set of resources as usable by a bounded set of agents. In other words, one of the advantages of A_1 may be not the ability to pursue a given project p_1 with r_2 better than A_2 could have, but to see that a more valuable p_2 is possible. Second, the initial statement does not take into consideration collaboration, and the possible ways in which cooperating individuals can make each other creative in different ways than they otherwise might have been. Once one takes into consideration these diverse effects on the increased possibilities for relationships among individuals and between

individuals and resources, it becomes likely that there are increasing returns to increases in the number of agents and resources involved in a production process.

Assume, for example, that every agent, given a t value, has some potential to being able to use every resource, which could be measured as an option for that agent on that resource. (In other words, its value is derived from the value of the agent using the resource, discounted by the probability that the agent will be good at using the resource.) Assume also that every agent has some added potential to add value in collaboration with any other agent, and that every resource could have some potential value in combination with any other resource. If we have one agent A_1 and one resource r_1 , the value of the option of A_1 to use r_1 is the value of the resource set. If we add one more resource, we get the value of A_1 to use r_1 , A_1 to use r_2 , and A_1 to use r_1 in combination with r_2 . Symmetrically, if we keep the resource set fixed at one resource, but add an agent, because the resources are nonrival we would see the value of two agents and one resource as the sum of the values of A_1 to use r_1 , A_2 to use r_1 , and a collaboration between A_1 and A_2 to use r_1 . If we combine adding one agent and one resource, we see the following. The value of $\{A_1, r_1\} + \{A_2, r_2\}$, if the two sets are strictly separated, is the value of A_1 using r_1 and A_2 using r_2 . The value of $\{A_1, A_2, r_1, r_2\}$ in a single agent/resource space is the combined value of A_1 using r_1 and A_2 using r_2 , A_1 using r_2 and A_2 using r_1 , A_1 using r_1 and r_2 and A_2 using r_1 and r_2 , A_1 and A_2 collaborating to use r_1 , A_1 and A_2 collaborating to use r_2 , A_1 and A_2 collaborating to use r_1 and r_2 .⁹¹ Figures 3 and 4 illustrate this point. Every arrow identifies one potential option for a valuable combination of agents and resources. In Figure 3 we see that separating the two agents and resource results in a combined value of only two options. In Figure 4 we see that, once both agents and resources are placed in the same opportunity set, the number of options for use and collaboration increase dramatically, and each arrow represents one of the nine potentially valuable combinations of agents and resources that the combination makes possible.

⁹¹ I am not sure there is room to formalize the precise relationship here on the style of Metcalfe's Law or Reed's Law, see David P. Reed, That Sneaky Exponential—Beyond Metcalfe's Law to the Power of Community Building, <http://www.reed.com/Papers/GFN/reedslaw.html>. From a policy perspective, there is no need to do so at this early stage of studying the phenomenon. It is sufficient for our purposes here to see that the collaboration effects and insights due to exposure to additional resources mean that the returns to scale are, as with other networks, more than proportional.



If this is true, then in principle a state in which all agents can act effectively on all resources will be substantially more productive in creating information goods⁹² than a world in which firms divide the universe of agents and resources into bounded sets. As peer production relies on opening up access to resources for a relatively unbounded set of agents, freeing them to define and pursue an unbounded set of projects that are the best outcome of combining a particular individual or set of individuals with a particular set of resources, this open set of agents is likely to be more productive than the same set could have been if divided into bounded sets in firms.

This is not to say that peer production will always necessarily be more productive and that it will always improve with size. First, the observation is limited by the extent to which adding agents increases the coordination and communication costs and increases the probability that the set of agents will include agents whose actions, through incompetence or malice, will undermine the productivity of the set.

Second, where focused but relatively standardized effort is more important than variability of talent, well-understood incentive systems based on monetary rewards could outweigh this effect, and markets and firms are much better understood mechanisms to generate the incentives for such application of effort than is peer production. This effort effect, however, may be of limited effect, and may indeed even have the opposite effect. If a project can be structured to resolve the effort/incentives problem without appropriation of the output, the substantial increases in productivity born of the availability of a larger set of resources to a larger set of agents with widely variable talent endowments could be enough to make even an imperfectly motivated peer production process more productive than firms that more

⁹² Note that the effect changes dramatically when the resources are rival, because then the value of any agent or combination of agents working on the resource is not additive to the value of any other agent or combination, since the use by one excludes the use by others in a way that is not true for a purely non-rival good like information. The allocation gain is attained in allocating the scarce resource—human attention, talent, and effort—given the presence of non-rival resources to which the scarce resource is applied with only a probability of productivity.

directly motivate effort but segment agents and resources into smaller bounded sets. Moreover, as Part III.A explains, a peer production project could increase, rather than decrease, motivation, by eliciting contributions motivated by non-monetary motivations when monetary rewards would have been either ineffective or inefficient.

And third, as unbounded sets of agents utilize unbounded sets of resources for unbounded sets of projects, there is likely to be substantial duplication of effort. Plainly this duplication is wasteful if one considers actual likely patterns of peer production as compared to an idealized peer production system where everyone self-identifies perfectly for the one contribution that they are best suited to produce. But this plainly will not happen. The question then becomes one of comparative efficiency—how much of a drag duplication is on the claimed increased efficiency provided by peer production enterprises. The answer has two primary components. First, as Part III will elaborate, peer production draws effort that in many cases would otherwise have been use in purely non-productive consumption—say, watching television instead of marking craters on Mars, ranking websites for the Open Directory Project, or authoring entries for Wikipedia. On a macro level of social productivity, then, an economic system that incorporates peer production as one component in its production system will add a vehicle for tapping effort pools that would otherwise not be used productively at all. While duplication may limit the total value of this newly tapped source of productivity, it is less important to the extent that the duplication occurs among efforts that would, in the absence of a peer production system, have gone unused in the production system. Second, and probably more important, redundancy provides important values in terms of the robustness and innovativeness of a system. Having different people produce the same component makes the production system more robust to occasional failures. Moreover, having different people with different experience and/or creative approaches attack the same problem will likely lead to an evolutionary model of innovation, where alternative solutions present themselves, giving the peer production process the ability to select among a variety of actual solutions rather than pre-committing to a single solution.

III. Of Motivation and Organization: The Commons Problem

1. The “incentives” problem: of diverse motivations and small contributions

What makes contributors to peer production enterprises tick? Why do they contribute? There are two versions of this question. The first is the question of the economic skeptic. It questions the long term sustainability of this phenomenon, given that people will not, after the novelty wears off, continue to work on projects in which

they can claim no proprietary rights.⁹³ It is to this question that my discussion here responds, in an effort to show that the network as a whole can be a sustainable system for the production of information and culture. There is a second, narrower version of the question, which arises once one overcomes the skepticism and begins to consider how peer production can be steered or predicted. It would seek to understand the motivations and patterns of clustering around projects in the absence of property rights and contracts, and the emergence of the effective networks of peers necessary to make a *particular* project succeed. These are questions that present rich grounds both for theoretical and empirical study. My hunch is that these would best be done informally in the domains of social psychology and anthropology, or, if done formally, through artificial life-type modeling. They are, in any event, beyond the scope of this initial study, which is intended solely to define the phenomenon and assess its sustainability and welfare effects in general terms.

The incentive problem as an objection to the general sustainability of peer production is in large part, as a practical matter, resolved by the existence of a series of mechanisms for indirect appropriation of the benefits of participation catalogued quite comprehensively by Lerner & Tirole.⁹⁴ At the broadest level, there is the pleasure of creation. Call it dispassionately “hedonic gain” or romantically “an urge to create,” the mechanism is simple. People are creative beings. They will play at creation if given an opportunity, and the network and free access to information resources provide this opportunity.⁹⁵ More closely related to the project of keeping

⁹³ This skepticism is more often encountered in questions in conferences and presentations than in formal papers. A well articulated written example of a skeptic’s view, however, is R. L. Glass, The sociology of open source: of cults and cultures. *IEEE Software*, May/Jun (2000) (Comparing recruiting OS developers to Tom Sawyer’s whitewashing the fence trick, arguing that eventually OS efforts will die because too many important programming tasks are not fun/sexy enough).

⁹⁴ Lerner & Tirole, Simple economics of open source, *supra*.

⁹⁵ Moglen makes this central to his explanation. See Moglen, Anarchism Triumphant, *supra* note ____ (“It’s an emergent property of connected human minds that they create things for one another’s pleasure and to conquer their uneasy sense of being too alone. The only question to ask is, what’s the resistance of the network? Moglen’s Metaphorical Corollary to Ohm’s Law states that the resistance of the network is directly proportional to the field strength of the ‘intellectual property’ system.” and also, “So, in the end, my dwarvish friends, it’s just a human thing. Rather like why Figaro sings, why Mozart wrote the music for him to sing to, and why we all make up new words: Because we can. Homo ludens, meet Homo faber.”). Raymond, Homesteading the Noosphere, 13-14; and Lerner & Tirole *supra*, also offer hedonic gains as one component of their respective explanations. There is, of course, something counter-intuitive about calling hedonic pleasure “indirect” appropriation. I use the terms “direct” and “indirect” to distinguish between appropriation that relies directly on the economic exclusion made possible by intellectual property law, and all other forms of appropriation. The distinction is made so as to enable us to map whenever talking about appropriation whether it is a form of appropriation that supports the utility of intellectual property rights or a form of alternative appropriation, which undermines the justification of intellectual property rights. On why it is that these two types of appropriation have this oppositional relationship as insights into the utility of the intellectual property rights see Yochai Benkler, Intellectual Property and the Organization of Information Production, 22 Int’l Rev. L. & Ec. 81 (2002) (explaining

body and soul together, there are a variety of indirect appropriation mechanisms for those who engage in free software development. These range from the amorphous category of reputation gains,⁹⁶ through much more mundane benefits such as consulting contracts, customization services, and increases in human capital that are paid for by employers who can use the skills gained from participation in free software development in proprietary projects. In this regard, it is important to note that about two thirds of the revenues in the software industry are not tied to software publishing, but to service-type relationships.⁹⁷ Given that the software industry is more than three times the size of, for example, movie, video, and sound recording put together,⁹⁸ two thirds of this industry still provides a field from which to extract value for software developers who participate in free software projects that is more than twice the size of those smaller, albeit more flamboyant industries.

The reality of phenomena like academic research, free software, the World Wide Web, NASA's clickworkers or Slashdot supports these explanations with robust, if not quantified here, empirical grounding. All one need do is look at the Red Hat founders (no longer billionaires, but not quite on the bread line, either)⁹⁹ and IBM's billion-dollar commitments to supporting Linux and Apache on the one hand,

that intellectual property has positive effects on information production strategies that are based on direct appropriation and negative effects on information production strategies that rely on indirect appropriation, and that these effects will structure the organization of information production and its efficiency even when they have no affect on aggregate productivity in information production).

⁹⁶ But, as mentioned above, this commonly cited motivation has not been reconciled with the contrary practices of two of the most successful free software projects, Apache and the Free Software Foundation, neither of which provides personal attribution to code they bless. See *supra*, note 7.

⁹⁷ The Economic Census of 1997 breaks up software into several categories, ranging from publishing to different types of services and education. Nonetheless, it is possible to collect the information about the industry as a whole using primarily the following categories: software publishing (32.4%), computer systems design and related services (57.9%), systems consultants, (8.5%), and computer training (1.2%). Software publishing (NAICS 5112) (the things we usually think about when we think about software) had receipts of over 61 billion dollars; Computer systems design and related services (NAICS 5415) (defined as "This industry comprises establishments primarily engaged in providing expertise in the field of information technologies through one or more of the following activities: (1) writing, modifying, testing, and supporting software to meet the needs of a particular customer; (2) planning and designing computer systems that integrate computer hardware, software, and communication technologies; (3) on-site management and operation of clients' computer systems and/or data processing facilities; and (4) other professional and technical computer-related advice and services," roughly 109 billion dollars; systems consultants (NAICS 5415122) roughly 16 billion dollars, and computer training, roughly 2.5 billion dollars.

⁹⁸ See 1997 Economic Census. All movie, video, and recording industries (NAICS 512) had total receipts of roughly 56 billion dollars, as compared to roughly 188 billion dollars for the software industry, see *supra*, notes 97.

⁹⁹ Red Hat is a company that specializes in packaging and servicing GNU/Linux operating systems. In 1999 it had an immensely successful IPO that made its founders billionaires, for a while. The company survived the Bubble bursting and continues to lead the field of Linux distributions.

and the tens of thousands of volunteer clickworkers, thousands of Linux developers, and hundreds of distributed proofreaders, on the other hand, to accept intuitively that some combination of hedonic gain and indirect appropriation can resolve the incentives problem. In this part I abstract from this intuitive observation to offer an answer that is more analytically tractable and usable to understand the micro-analytic questions of peer production and the potential range in which peer production will be more productive than firms or markets.

a. Abstracting the effect of diverse motivations

Saying that people participate for all sorts of reasons is obviously true at an intuitive level. It does not, however, go very far towards providing a basis for understanding why some projects draw many people, while others fail, or how the presence or absence of money affects the dynamic. What I will try to do in this section is to propose a framework to generalize the conditions under which peer production processes will better motivate human effort than market-based enterprises. Given the discussion of the information and allocation gains offered by peer production, this section outlines a range in which peer production should be more productive than market-based or firm-based production. At the broadest level, wherever peer production can motivate behavior *better* than markets or firms, then certainly it will be superior. It will also be potentially better over a range where it may motivate behavior less effectively than markets or firms, but the contribution of the lower overall effort level will be less than the contribution of the added value in terms of information about and allocation of human creativity.

Let any agent have a set of preferences for rewards of three types:

- M => *Monetary rewards*, which decrease in value because of the decreasing marginal utility of money. Call the rate at which M decreases s (satiation).
- H => *Intrinsic hedonic rewards* experienced from taking the actions
- SP => *Socio-psychological rewards*, which are a function of the cultural meaning associated with the act, and may take the form of actual effect on social associations and status perception by others, or on internal satisfaction from one's social relations or the culturally-determined meaning of one's action, etc.

At an intuitive level, three common examples help to clarify this diversity of motivation. Simplest to see is how these motivations play out with regard to sex: the prostitution fee (M), the orgasm (H), and love (SP). One can also make and serve dinner to others for any combination of a fee, the pleasure of cooking, and companionship, and which combination of these is involved in an interaction will

shape our understanding of whether we are observing a short-order cook, a restaurant chef, or a dinner party host. Similarly, one can write about law for a legal fee, the pleasure of creating a well-crafted argument, or the respect of the legal community or one's colleagues. To some extent, all three exist for anyone writing, but in different measures depending on social role, such as whether the author is a practitioner, a judge, or an academic, or on other factors, such as how time-constrained they are.

The value of the three types of rewards for any given action might be independent of the value of the others, or it might not.¹⁰⁰ For purposes of this analysis I will assume that H is a personal preference that is independent of the other two,¹⁰¹ but that M and SP can be positively or negatively correlated depending on the social construction of having money associated with the activity. I will call this factor *p*, which can be negative (as in prostitution) or positive (as in professional sports).¹⁰² The *p* factor is most interesting when it is negative, and is intended to allow for the possibility of a "crowding-out" phenomenon,¹⁰³ which has mostly been studied in the context of the relatively rare instances where altruistic provisioning has been the major, if not exclusive, mode of provisioning of socially important material goods, at least in some societies, such as blood¹⁰⁴ or gametes.¹⁰⁵ While analysis leaves serious

¹⁰⁰ Needless to say, the independent value of each may be positive or negative. One might be willing to pay money to engage in a hedonically pleasing or socio-psychologically satisfying activities, as people do all the time for hobbies, and people often take hedonically unpleasant or socially awkward or even demeaning jobs in order to get the positive monetary rewards.

¹⁰¹ Though separating out purely physical pleasure or pain from the social-psychological meaning of the cause of the pleasure or pain is artificial in the extreme. In principle, hedonic gain can be treated as part of SP, and indeed I ignore it as an independent factor in the analysis. I have it in the general statement largely to separate out the social-psychological aspect, which, unlike hedonic gains, is usually submerged in economics.

¹⁰² Again, the culturally contingent nature of the relationship should be obvious. When the Olympics were renewed in the modern era, they were limited to "amateurs" because professional sports were a form of entertainment, giving their paid performers no more respect than paid performers were given more generally. As with all performers, this changed with the status inversion that was part of the 20th century celebrity culture generated to focus mass demand on mass-produced entertainment, as opposed to the relationship/presence-based entertainment of the past.

¹⁰³ See Bruno S. Frey and Reto Jege, Motivation Crowding Theory: A Survey of Empirical Evidence, 15(5) J. Economic Surveys 589 (2001); Bruno S. Frey and Felix Oberholzer-Gee, The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out, 87 *Am Econ. Rev.* 746 (1997); see also Fehr and Gächter, *supra*. For a broader moral claim about this tradeoff see Margaret Jane Radin, Contested Commodities, (1996), and a critique in Kenneth J. Arrow, Invaluable Goods. 35 J. Econ. Lit. 757 (1997).

¹⁰⁴ The quintessential source of the claim that altruism is superior to markets in providing blood is Richard M. Titmuss, *The Gift Relationship; From Human Blood to Social Policy* (1972).

¹⁰⁵ See L. S., Fidell *et al.* Paternity by proxy: Artificial insemination with donor sperm, in *Gender in Transition: A New Frontier*, (J. Zuckerberg ed.), at 93 (1989) (reporting that three-quarters of sperm donors in the USA were primarily motivated by financial gain); K. R. Daniels, Semen donors: Their motivations and attitudes to their offspring, 7 *J. Reprod. Infant. Psychol.* 121 (1989) (finding that "a majority of donors surveyed in Australia and New Zealand gave altruistic reasons as motivation, with payment rated as rather unimportant"); S. B. Novaes, Semen Banking and artificial insemination by

questions as to whether altruistic provisioning of these types of goods is indeed superior to market-based provisioning as a general social policy,¹⁰⁶ the primary disagreement concerns which is more efficient in the aggregate, not whether market provisioning displaces altruistic provisioning and whether each mode draws different contributors.¹⁰⁷ Using our three intuitive examples, an act of love drastically changes meaning when one person offers the other money at its end, and a dinner party guest who will take out a checkbook at the end of dinner instead of bringing flowers or a bottle of wine at the beginning will likely never be invited again. The question of money in legal writing will depend on role. It will have a different effect if the social construction of the role of the author. For a practicing advocate, p usually is positive, and higher monetary rewards represent the respect the author receives for her craft. For a judge, p with regard to payment for any particular piece of writing is strongly negative, representing the prohibition on bribes. For academics, p for a particular piece of writing may be positive or negative, depending on whether its source is considered to be an interested party paying for something that is more akin to a brief than an academic analysis, or, for example, a foundation or a peer-reviewed grant, in which case “winning” the support is considered as adding prestige.

A distinct motivational effect arises when SP is associated with participation in collective action, and concerns the presence or absence of rewards to the other participants and the pattern of the reward function—that is, whether some people get paid and others do not, or if people get paid differentially for participating. This relationship could be positive where altruism or a robust theory of desert culturally structures the social-psychological component of the reward to support monetary

donor in France: Social and Medical Discourse, 2 *Int. J. Technol. Assessm. Health Care* 219 (1986) (providing an account of the several sperm banks in France and an analysis of their varying policies on donor compensation); S. B. Novaes, Giving, receiving, repaying: Gamete donors and donor policies in reproductive medicine, 5 *Int. J. Technol. Assessm. Health Care* 639 (1989) (reviewing motivation and social issues of sperm and egg donation and surrogacy).

¹⁰⁶ Titmuss's thesis was challenged in a series of papers in the 1970s, see, e.g., Kenneth J. Arrow, Gifts and Exchanges, 1 *Philosophy and Public Affairs* 343 (1972); Robert S. Solow, Blood and Thunder, 80 *Yale L. J.* 1696 (1971), and more recently has been subject to refinement with the experience of the AIDS epidemic, see Kieran Healy, The Emergence of HIV in the U.S. Blood Supply: Organizations, Obligations, and Management of Uncertainty, 28 *Theory and Society* 529 (1999).

¹⁰⁷ Specifically for an evaluation of Titmuss's argument in light of the HIV crisis see, e.g., Kieran Healy, Embedded Altruism: Blood Collection Regimes and the European Union's Donor Population, 105 *American Journal of Sociology* 1633 (2000) (reporting on an international comparison and concluding that “the opportunity to sell plasma does reduce one's likelihood of giving blood”). More generally, for a description of empirical surveys in a number of areas, see Frey and Jegen, *supra* note 101 (describing empirical research in multiple disciplines supporting the displacement effect money has on voluntaristic motivations).

appropriation by others,¹⁰⁸ or, more commonly perhaps, negative, where one agent is jealous of the rewards of another.¹⁰⁹ I will denote this factor *jalt* (jealousy/altruism).

Agents will then face different courses of action that they will perceive as having different expected rewards R:

$$R = M_s + H + SP_{p,jalt}$$

At any given time, an agent will face a set of possible courses of action, and will have a set of beliefs about the rewards to each course of action, each with this form. A rational agent will choose based on the value of R, not of M. Irrespective of one's view of whether the agent is a maximizer or a satisficer, the agent will have some total valuation of the rewards to differing courses of action, and hence of the opportunity cost of following courses of action that exclude other courses of action.

It is quite intuitive to see then that there will be some courses of action whose reward will be heavily based on hedonic or social-psychological parameters, a combination of all three factors, or primarily monetary rewards. At the broadest level one can simply say that agents will take actions that have a positive value and low opportunity cost because they do not displace more rewarding activities. Similarly, where opportunities for action do compete with each other, an agent will pursue an activity that has low, no, or even negative monetary rewards when the total reward, given the hedonic and social-psychological rewards, is higher than alternative courses of action that do have positive monetary rewards attached to them. Hence the phenomena of starving artists who believe they are remaining true to their art rather than commercializing, or of law professors who forego large law-firm partner draws when they choose teaching and writing over the practice of law.

What more can we say about the likely actions of agents whose preferences for rewards have the form I describe? First, there is a category of courses of action that will only be followed, if at all, by people who seek social-psychological and hedonic rewards. Assume that there are transaction costs for defining and making M and SP available to the agent, C_m and C_{sp} , respectively. I assume that these costs are different, because the former require definition and enforcement of property rights, contracts, and pricing mechanisms, while the latter require social mechanisms for the association of social-psychological meaning with the act generically and with the individual agent's act in particular.

¹⁰⁸ A religiously motivated agent, for example, might consider the acquisition of monetary returns by other agents a positive sign of success, because the appropriators are seen as deserving—whatever theory of desert is prescribed by the religion—say, neediness or having been chosen in some sense.

¹⁰⁹ Genesis 4:3-8.

There is potentially a category of cases where the marginal value (V) of an agent's action will be less than the transaction costs of providing monetary rewards for it, in which case the expected monetary reward will be zero. If the social value of the contribution is greater than zero, however, and if the hedonic and social-psychological rewards are greater than zero and greater than the cost of making the social-psychological rewards available, then it will be socially efficient for agents to act in this way when opportunities to act arise. Agents will in fact do so if someone has incurred the costs of providing the opportunities for action and the social-psychological or hedonic rewards.

Behaviors in the following range will therefore occur only if they can be organized in a form that does not require monetary incentives and captures behaviors motivated by social-psychological and hedonic rewards:

$$C_m > V > C_{sp}, \text{ and } SP - C_{sp} + H > 0$$

Whether this range of activities is important depends on the granularity of useful actions. The more fine-grained the actions, and the more of these small-scale actions need to be combined into a usable product, the higher the transaction costs of monetizing them relative to the marginal contribution of each action.¹¹⁰

Second, approaches that rely on social-psychological rewards will be particularly valuable to motivate actions that are systematically undervalued in the market, because they generate high positive externalities. A fairly intuitive example is basic science, which is particularly ill suited for proprietary information production because of its high positive externalities,¹¹¹ and where our social-cultural framework has developed an elaborate honor-based rewards system rather than one focused on monetary rewards. We see similar social-psychological reward structures to reward and motivate participation in other practices that produce high-positive externalities that would be difficult fully to compensate in monetary terms, like teaching, military service or uncorrupt political, cultural, or spiritual leadership. Similarly, to the extent

¹¹⁰ If these transaction costs can be lowered by technology, this would counteract the effect and decrease the size of the group of cases that fall into this category.

¹¹¹ The public goods problem of information production limits the efficacy of proprietary provisioning under any circumstances. The fact that basic science has many and varied uses as a fundamental input into new innovation and learning, that is, creates particularly large positive externalities, makes the public goods problem particularly salient in that context. See Richard R. Nelson, *The Simple Economics of Basic Scientific Research*, 48 *Journal of Political Economy* 297-306 (June 1959); Arrow, *Economic Welfare and the Allocation of Resources for Invention*, *supra*, at 623-25; Eisenberg, *cDNA sequencing*, *supra*; Nelson, *What is "Commercial" and What is "Public" about Technology, and What Should Be?* in *Technology and the Wealth of Nations* at 65-70 (Rosenberg, Landau, and Mowery 1992); Ralph Gomory, *The Technology-Product Relationship: Early and Late Stages*, in Rosenberg, Landau, and Mowery, *supra*, at 388.

that peer production can harness motivations that do not require monetization of the contribution, the information produced using this model can be released freely, avoiding the inefficiencies associated with the public goods problem of information.

It is important to recognize that actions involved in creating the opportunities for others to act are themselves acts with similar reward structures. The scientists who created the Mars clickworkers project operated on one set of monetary and social-psychological returns, while the clickworkers themselves operated in response to a different set of hedonic and social-psychological returns. The Open Source Development Network funds the Slashdot platform based on one set of rewards, that includes an expected monetary return, but its action generates opportunities for others to act purely on SP and H type rewards. The crucial point is that the presence of M type rewards for the agent generating the opportunities does not negatively affect the social-psychological returns to agents who act on these opportunities. In other words, that there be some reason why the different reward structure will not give the *jalt* factor for the contributors a strong negative value based on the monetary rewards captured by the person providing the opportunity for collaboration.

We need, then, to state the relationship between the presence of M type rewards for an action and the SP type rewards associated with it. For simplicity I will treat the total effect of both modifiers of SP as p , and will separate out *jalt* only where there is a reason to differentiate between effect of monetary returns to the agent and effects of differential reward functions for different agents in a collaborative group—as in the case where the person offering the opportunity to collaborate has different rewards from the participants in the collaboration.

Keeping hedonic gains to one side, the reward function looks like this.

$$R = M_s + SP_p$$

First, we can confidently say that whenever M and SP are independent of each other or are positively correlated (that is, when $p \geq 0$), approaches that provide monetary rewards for an activity will dominate non-monetizable approaches towards the exact same activity. A rational agent will prefer a project that provides both social-psychological and monetary rewards than one that offers only one of these rewards. Someone who loves to play basketball will, all other things being equal, prefer to be paid for playing at Madison Square Garden over playing at West 3d and Sixth Avenue without being paid.

Third, we can say that when M and SP are negatively correlated ($p < 0$), an activity will be more or less attractive to agents depending on the values of s and p , that is, on the rate at which the value of marginal monetary rewards for a new action is

discounted by the agent and the rate at which the presence of money in the transaction devalues the social-psychological reward for that action. Table 3 maps the likely effects of monetary rewards on the value of R as a function of the values of s and p . We can say generally that individuals with a high discount rate on money (high s) will be likely to pursue activities with a high absolute value negative p rate *only* if these are organized in a non-proprietary model, because the value of M_s for them is low, while the presence of any M-type reward substantially lowers the value of SP_p . At the most simple level, this could describe relatively wealthy people—for example, a wealthy person is unlikely to take a paying job serving lunch at a soup kitchen, but may volunteer for the same job. More generally, most people who have finished their day job and are in a day part that they have chosen to treat as leisure, even though a second job is available, can be treated as having a higher s value for that day part. In that day part, it will likely be easier to attract people to a project with socio-psychological benefits, and if p is large and negative, adding monetary rewards will lower, rather than increase, participation. As we move towards a situation where the value of s for an individual is low, and the p rate, though negative, is low, we will tend to see a preference for combining M and SP , as one would where p is neutral or positive.

$\begin{matrix} * \\ * \\ p < 0 \\ * \\ s \\ * \end{matrix}$	High	Low
	High	Low
High	Monetary rewards lower R substantially	Monetary rewards decrease R through SP , but increase R through M , total R may be improved
Low	Monetary rewards affect R only slightly in either direction	Monetary rewards likely increase R

Table 3: Effects of monetary rewards on the total rewards seen to an activity as a function of s and p .

Finally, there may be ways in which p can be changed from negative to positive, or its negative value can be reduced, by changing the way M is correlated to the action. To stay with the sex example, while there is some social discomfort associated with marriage “for money”, it does not approach the level of social approbation associated with prostitution. The p value is negative, but smaller. In other societies, perhaps in times holding less egalitarian ideals about marriage, there might have actually been a positive p value—as in “a good catch.” Similarly,

professional performers or athletes may have been treated with less respect than amateurs a hundred years ago, but this has obviously changed quite dramatically. The same can be said for the *jalt* factor. One can imagine that free software development communities would attach a negative social value to contributions of those who demand to be paid for their contributions. The same communities may have different feelings towards programmers who contributed for free, but who later get large consulting contracts as a result of the experience and reputation they gained from their freely shared contributions.

This analysis suggests a series of likely conditions under which nonproprietary organizational approaches will be sustainable. First, there is the case of projects that are broken down into fine-grained modules, where market remuneration would likely be too costly to sustain, but where hedonic and social-psychological rewards can provide contributors with positive rewards. As I will explain in the following section, fine-grained modularity is an important characteristic of the large-scale collaborations that form the basis of peer production. The analysis of motivations suggests that peer production will not likely be effectively harnessed using direct monetary incentives. Second, there are instances where the value of monetary return is small relative to the value of the hedonic and social-psychological rewards, particularly where the cultural construction of the social-psychological rewards places a high negative value on the direct association of monetary rewards with the activities. Teenagers and young adults with few economic commitments and a long time horizon for earning and saving, on the one hand, and high social recognition needs, on the other hand, are an obvious group fitting these characteristics. Another group comprises of individuals who have earnings sufficient to serve their present and expected tastes, but who have a strong taste for additional hedonic and social-psychological benefits that they could not obtain by extending their monetarily remunerated actions. Academics in general, and professional school academics in particular, are obvious instances of this group. Many of the volunteers to Internet-based projects who volunteer instead of watching television or reading a book likely fall into this category. Individuals whose present needs are met but whose future expected needs require increased monetary returns might participate if the social-psychological returns were not negatively correlated with future, indirect appropriation, such as through reputation gains. This would effectively mean that they do add an *M* factor into their valuation of the rewards, but they do so in a way that does not negatively affect the value of *SP* for themselves or for other contributors to collaborative projects.

b. Diverse motivations and large-scale collaborations

The diversity of motivations allows large-scale collaborations to convert the motivation problem into a collaboration problem. In other words, the motivation problem is simple to resolve if the efforts of enough people can be pooled.

In a corollary to “Linus’s Law,”¹¹² one might say

Given a sufficiently large number of contributions, direct monetary incentives necessary to bring about contributions are trivial.

The “sufficiently large” aspect of this observation requires some elaboration. “Sufficiently” refers to the fact that the number of people who need to collaborate to render the incentives problem trivial depends on the total cost or complexity of a project. The sustainability of any given project depends, however, not on the total cost but on how many individuals can contribute to it relative to the overall cost. If a project that requires thousands of person hours can draw on the talents of 30,000 or 15,000 individuals instead of a few dozen or a few hundred, then the contribution of each, and hence the personal cost of participation that needs to be covered by diverse motivations, is quite low. Similarly, a project that requires ten or twenty person hours can be provided with little heed to incentives if it can harness the distributed efforts of dozens of participants.

More generally, one can state:

Peer production is limited not by the total cost or complexity of a project, but by its modularity, granularity, and the cost of integration.

Modularity is a property of a project referring to the extent to which it can be broken down into smaller components, or modules, that can be independently and asynchronously produced before they are assembled into a whole. If modules are independent, individual contributors can choose what and when to contribute independently of each other, thereby maximizing their autonomy and flexibility to define the nature, extent, and timing of their participation in the project. Given the centrality of self-direction of human creative effort to the efficiencies of peer production, this characteristic is salient.

Granularity refers to the size of the modules, in terms of the time and effort that an agent must invest in producing them. The number of people who will likely

¹¹² Coined by Eric Raymond to capture one of the attributes of the approach that developed Linux: “Given enough eyeballs, all bugs are shallow.” Raymond, Cathedral and Bazaar, *supra*, at 9.

participate in a project is inversely related to the size of the smallest-scale contribution necessary to produce a usable module. Usability may place a lower boundary on granularity either for technical or economic reasons, where at a minimum the cost of integrating a component into a larger modular project must be lower than the value that adding that component adds to the project. But above that boundary, the granularity of the modules sets the smallest possible individual investment necessary to participate in a project. If this investment is sufficiently low, then incentives for producing that component of a modular project can be of trivial magnitude and many people can contribute. If the finest-grained contributions are relatively large and would require large investment of time and effort, the universe of potential contributors decreases. A successful large-scale peer production project must therefore have a predominate portion of its modules be relatively fine-grained. The discussion in the preceding section suggests that, given the relatively small independent value such fine-grained contributions will have and the transaction costs associated with remunerating each contribution monetarily, non-monetary reward structures are likely to be more effective to motivate peer production efforts.

Independent of the minimal granularity of a project, heterogeneity in the size of the modules may add to its efficiency. Heterogeneity allows contributors with diverse levels of motivation to collaborate by contributing modules of different sizes, whose production therefore requires different levels of motivation. Contributors may vary widely in their hedonic taste for creation, their social-psychological attitude towards participation, or in opportunities for indirect monetary appropriation (like the difference between IBM or Red Hat and individual volunteers in free software projects). A project that allows highly motivated contributors to carry a heavier load will be able to harness a diversely motivated human capital force more effectively than a project that can receive only standard-sized contributions.

2. *Integration: Problem and opportunity*

The remaining obstacle to effective peer production is the problem of *integration* of the modules into a finished product. Integration includes two distinct components—first, a mechanism for providing quality control or integrity assurance, to defend the project against incompetent or malicious contributions, and second, a mechanism for combining the contributed modules into a whole. It is here that the term “commons” that I use in describing the phenomenon as “commons-based peer production” gets its bite, denoting the centrality of the *absence of exclusion* as the organizing feature of this new mode of production, and highlighting the potential pitfalls of such an absence for decentralized production. Observing commons-based peer production on the background of the commons literature, we see integration and the commons problem it represents solved in peer production efforts by a combination of four mechanisms: iterative peer production of the integration function itself,

technical solutions embedded in the collaboration platform, norm-based social organization, and limited reintroduction of hierarchy or market to provide the integration function alone. In all events the integration function must be either low-cost or itself sufficiently modular to be peer-produced in an iterative process, in order for a project to be susceptible to sustainable peer production.

What kind of commons is it, then, that peer production of information relies upon? Commons are most importantly defined by two parameters.¹¹³ The first parameter is whether use of the resource is common to everyone in the world or to a well-defined subset. The term “commons” is better reserved for the former, while the latter is better identified as a common property regime (CPR)¹¹⁴ or limited common property regime.¹¹⁵ The second parameter is whether use of the resource by whoever the set of people whose use is privileged is regulated or not. Here one can more generally state, following Rose,¹¹⁶ that resources in general can be subject to regimes ranging from total (and inefficiently delineated) exclusion—the phenomenon Heller has called the anticommons¹¹⁷—through efficiently-delineated property and otherwise regulated access, to completely open, unregulated access. The infamous “tragedy of the commons” is best reserved to refer only to the case of unregulated access commons, whether true commons or CPRs. Regulated commons need not be tragic at all, and indeed have been sustained and shown to be efficient in many cases.¹¹⁸ The main difference is that CPRs are usually easier to monitor and regulate—using both formal law and social norms¹¹⁹—than true commons, hence the latter may more often slip into the open access category even when they are formally regulated.

¹¹³ The most extensive consideration of commons and the resolution of the collective action problems they pose is Ostrom, *supra*.

¹¹⁴ See Ostrom, *supra*.

¹¹⁵ Carol M. Rose, The Several Futures of Property: Of Cyberspace and Folk Tales, Emission Trades and Ecosystems, 83 Minn. L. Rev. 129 (1998).

¹¹⁶ Carol M. Rose, Left Brain, Right Brain and History in the New Law and Economics of Property, 79 Org. L. Rev. 479 (2000).

¹¹⁷ Michael A. Heller, The Tragedy of the Anticommons: Property in the Transition from Marx to Markets, 111 Harv. L. Rev. 621 (1998). I refer to Heller, rather than to Michelman, who to the best of my knowledge coined the term, see Frank I. Michelman, Ethics, Economics, and the Law of Property, in *Nomos XXIV: Ethics, Economics, and the Law 3* (J. Roland Pennock & John W. Chapman eds., 1982), because the concept, applied to inefficiently defined property rights relative to the efficient boundaries of resources as opposed to resources as to which everyone has a right to exclude, took off with Heller's use.

¹¹⁸ Ostrom, *Governing the Commons*, is the most comprehensive survey. Another seminal study was James M. Acheson, *The Lobster Gangs of Maine* (1988). A brief intellectual history of the study of common resource pools and common property regimes can be found in Charlotte Hess & Elinor Ostrom, *Artifacts, Facilities, And Content: Information as a Common-pool Resource*, (paper for the “Conference on the Public Domain,” Duke Law School, Durham, North Carolina, November 9-11, 2001).

¹¹⁹ The particular focus on social norms rather than formal regulation as central to the sustainability of common resource pool management solutions that are not based on property is Ellickson's, *supra*.

Ostrom also identified that one or both of two economic functions will be central to the potential failure or success of any given commons-based production system. The first is the question of provisioning, the second of allocation.¹²⁰ This may seem trivial, but it is important to keep the two problems separate, because if a particular resource is self-renewing if allocated properly, then institutions designed to assure provisioning would be irrelevant. Fishing and whaling are examples. In some cases, provisioning may be the primary issue. Ostrom describes various water districts that operate as common property regimes that illustrate well the differences between situations where allocation of a relatively stable (but scarce) water flow exists, on one hand, and where provisioning of a dam is the difficult task, after which water is abundant relative to demand.¹²¹ Obviously, some commons will require both.

Information production entails purely a provisioning problem. Because information is nonrival, once it is produced no allocation problem exists. Moreover, commons-based provisioning of information in a ubiquitously networked environment may present a more tractable problem than provisioning of physical matter, and shirking or free riding may not lead quite as directly to non-production. This is so for three reasons. First, the modularity of the projects allows redundant provisioning of “dropped” components to overcome occasional defections without threatening the whole. Second, a ubiquitously networked environment substantially increases the size of the pool of contributors. At first glance this should undermine peer production, because, generally speaking, the likelihood of free riding increases as the size of the pool increases and the probability of social-norms-based prevention of free riding declines.¹²² But as the size of the pool increases, the project can tolerate increasing levels of free riding, as long as the absolute number of contributors responding to some mix of motivations remains sufficiently large so that the aggregation of the efforts of those who do contribute, each at a level no higher than his or her level of motivation dictates, will be adequate to produce the good. As long as free riders do not affirmatively undermine production, but simply do not contribute, the willingness of contributors to contribute should depend on their perception of the likelihood of success given the number of contributors, not on the total number of users. Indeed, for contributors who seek indirect appropriation through means enhanced by widespread use of the joint product—like reputation or service contracts—a high degree of use of the end product, even by “free riders” who did not contribute to

¹²⁰ “Provisioning” refers to efforts aimed at producing a particular good that, but for the actions described by the term “provisioning” would not exist. “Allocating” refers to decisions about how a good that exists, but is scarce relative to demand for it, will be used most efficiently.

¹²¹ Ostrom, *Governing the Commons*, 69-88.

¹²² On the relationship between how small and closely knit a group is, and its capacity to use social norms to regulate behavior see Robert C. Ellickson, *Order Without Law* (1991). On the importance of social norms in regulating behavior in cyberspace see Lawrence Lessig, *Code and Other Laws of Cyberspace* (1999).

providing it, increases the expected payoff.¹²³ Third, the public goods nature of the product means that free riding does not affect the capacity of contributors to gain full use of their joint product, and does not degrade their utility from it. This permits contributors who contribute in expectation of the use value of the good to contribute without concern for free riding.

There are, however, types of defection that are likely to undermine provisioning by adversely affecting either (a) motivation to participate or (b) the efficacy of participation. The first type covers actions that reduce the value of participation, be it the intrinsic hedonic or social-psychological components, or the expected longer term extrinsic values, that is, the monetary rewards to reputation, human capital, etc. The second type relates mostly to potential failures of integration, due to an absence of an integration process, or due to poor quality contributions, for example.

Threats to Motivation

There are two kinds of actions that could reduce the intrinsic benefits of participation. First is the possibility that a behavior will affect the contributors' valuation of the intrinsic value of participation. Two primary sources of negative effect seem likely. The first is a failure of integration, so that the act of individual provisioning is seen as being wasted, rather than adding some value to the world. This assumes that contributors have a taste for contributing to a successful joint project. Where this is not the case—if integration is not a component of the intrinsic value of participation—then failure to integrate would not be significant. The World Wide Web is an example where it is quite possible that putting up a web site on a topic one cares about is sufficiently intrinsically valuable to the author, even without the sense of adding to the great library of the web, that integration is irrelevant to the considerations of many contributors.

The more important potential “defection” from commons-based peer production, is unilateral¹²⁴ appropriation. Unilateral appropriation could, but need not, take the form of commercialization of the common efforts for private benefit. More generally, appropriation could be any act where an individual contributor tries to make the common project reflect his or her values too much, thereby alienating other

¹²³ This attribute causes Steven Weber to describe free software production not only as “nonrival,” but as “anti-rival”—by which he means that increasing returns to widespread use mean that consumption by many not only does not reduce the value of a good, as in nonrival goods, but actually enhances it. See Weber, *supra* note 8, MS at 28-29.

¹²⁴ As opposed to collective, as in the conversion of some aspect of the commons to a common property regime where, for example, high quality or consistent contribution to the commons could become a criterion for membership.

participants from the product of their joint effort. The common storytelling enterprise called LambdaMOO, and the well-described crises that it went through with individuals who behaved in antisocial ways—like forcing female characters to “have sex” that they did not want to have in the story¹²⁵—is a form of appropriation—taking control over the production process so as to make the joint product serve one’s own goals. In LambdaMOO the participants set up a social structure for clearing common political will in response to this form of appropriation.¹²⁶ In the examples I have described in this paper, the explicit adherence to a norm of objectivity in Wikipedia,¹²⁷ and similar references in Kuro5hin to the norm of high-quality writing are clear examples of efforts to use social norms to regulate this type of defection by substantive, rather than commercial, appropriation. Similarly, some of the software-based constraints on moderation and commenting on Slashdot and other sites have the characteristic of preventing anyone from taking too large a role in shaping the direction of the common enterprise, in a way that would reduce the perceived benefits of participation to many others. For example, limiting moderators to moderating no more than five comments in any three-day period, or using trill filters to prevent users from posting too often, are technical constraints on permitting anyone to appropriate the common enterprise called Slashdot.

Another form of appropriation that could affect valuation of participation is simple commercialization for private gain. The primary concern is that commercialization by some participants or even by non-participants will create a sucker’s reward aspect to participation. This is the effect I introduced into the abstract statement of diverse motivations as the *jalt* factor—the effect of monetary rewards for others on the perceived value of participation. One example of such an effect may have occurred when the early discussion moderators on AOL boards—volunteers all—began to realize how that their contributions were effectively going to increase the value of the company, and left. There is, however, an immensely important counter example—to wit, the apparent imperviousness of free software production, our paradigm case, to this effect. Some contributors have made billions, while some of the leaders of major projects have earned nothing but honor.¹²⁸ Query, though, whether the pattern would have held if it were the primary leader of a project, such as Linus Torvalds, rather than people less central to the Linux kernel development process, who had made money explicitly by selling the GNU/Linux operating system as a product. It is, in any event, not implausible to imagine that individuals would be more willing to contribute their time and effort to NASA or a nonprofit enterprise than

¹²⁵ Larry Lessig, *Code and other Laws of Cyberspace* (2000).

¹²⁶ See Julian Dibbell, *A Rape in Cyberspace or How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database Into a Society*, *The Village Voice*, December 21, 1993, pages 36 through 42, text available <ftp://ftp.lambda.moo.mud.org/pub/MOO/papers/VillageVoice.txt>.

¹²⁷ See *supra*, text accompanying note ____.

¹²⁸ In this too peer production is similar to academic production, where scientists see their basic research used, very often by others, as the basis for great wealth in which they do not share.

to a debugging site set up by Microsoft. Whether this effect exists, how strong it is, and what are the characteristics of instances where it is or is not important is a valuable area for empirical research.

In addition to intrinsic value of participation, there is also an important component of motivation that relies on the use value of the joint project and on indirect appropriation based on continued access to the joint product—service contracts, human capital etc. For such projects, defection again may take the form of appropriation, in this case by exclusion of contributors from the use value of the end product. (Why academics, for example, are willing to accept the bizarre system in which they contribute to peer review journals for free, sometimes even paying a publication fee, and then have their institutions buy this work back from the printers at exorbitant rates remains a mystery.) In free software, the risk of defection through this kind of appropriation is deemed a central threat to the viability of the enterprise, and the GNU GPL is designed precisely to prevent one person from taking from the commons, appropriating the software, and excluding others from it.¹²⁹ This type of defection, on its face, looks like an allocation problem—one person is taking more than their fair share. But again, this is true only in a metaphoric sense. The good is still intrinsically a public good, and is physically available to be used by everyone. Law (intellectual property) may create this “allocation problem” in a misguided attempt to solve a perceived provisioning problem, but the real problem is effect on motivation to provision, not an actual scarcity that requires better allocation. The risk of this kind of unilateral appropriation lowers the expected value contributors can capture from their contribution, and hence lowers motivation to participate and provide the good.

Provisioning Integration

Another potential problem that commons based peer production faces is provisioning of the integration function itself. It is important to understand from the discussion here that integration requires some process for assuring the quality of individual contributions. This could take the form of (a) hierarchically managed review, as in the Linux kernel or Apache development processes, (b) peer review, as in the process for moderating Slashdot comments, (c) norm-based social organization, as in Wikipedia's objectivity norm, or (d) aggregation and averaging of redundant contributions, as in the Mars Clickworkers project. Academic peer production of science is traditionally some combination of the first three, although the Los Alamos

¹²⁹ The General Public License, under which GNU/Linux is distributed, is available at <http://www.fsf.org/copyleft/gpl.html>, section 2(b) limits the license to modify software distributed under the GPL such that the licensee “must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.”

Archive¹³⁰ and the Varmus proposal for changing the model of publication in the health and biomedical sciences¹³¹ towards free online publication coupled with post-publication peer commentary as a check on quality would tend to push the process further towards pure peer review and norms-based enforcement of the core values of completeness and accuracy, as well as attribution and respect for priority.

The first thing to see from the discussion of threats to motivation is that provisioning integration by permitting the integrator to be the residual owner (in effect, to “hire” the contributors and act as the entrepreneur) presents substantial problems for the motivation to provision in a peer-based production model. Appropriation may so affect motivation to participate that the residual owner will have to resort to market- and hierarchy-based organization of the whole production effort. Second, property rights in information are always in some measure inefficient.¹³² Creating full property rights in any single agent whose contribution is only a fraction of the overall investment in the product is even less justifiable than doing so for a person who pays all of the production costs. Third, and related, integration is quite possibly, particularly with the introduction of software-based management of the communications and to some extent the integration of effort, a low-cost activity. To the extent that this is so, even though integration may require some hierarchy, or some market-based provisioning, it is a function that can nonetheless be sustained on low returns—be it by volunteers, like those who run integration of code into the Linux kernel, by publicly-funded actors, as in the case of NASA clickworkers, or by firms, like the Open Source Development Network that supports Slashdot, that appropriate the integration value they add by means less protected from competition than intellectual property rights-based business models (in that case, advertising and service contracts to implement the same platform within business organizations).

The cost of integration—and hence the extent to which it is a limit on the prevalence of peer production—can be substantially reduced by both automation and the introduction of an iterative process of peer production of integration itself. First, integration could be a relatively automated process for some products. NASA clickworkers’ use of automated collation of markings and averaging out of deviations is an example, as are many of the attributes of Slashdot or Kuro5hin. Second, the integration function itself can be peer produced. Again with Slashdot, the software that provides important integration functions is itself an open-source project—in other words, peer-produced. The peer review of the peer reviewers—the moderators—is again distributed, so that 90% of registered users can review the moderators, who in

¹³⁰ <http://www.ArXiv.org>.

¹³¹ Harold Varmus, E-BIOMED: A Proposal for Electronic Publications in the Biomedical Sciences (1999) <http://www.nih.gov/about/director/pubmedcentral/ebiomedarch.htm>

¹³² See *supra*, note 73.

turn review the contributors. As peer production is iteratively introduced to solve a greater portion of the integration function, the residual investment in integration that might require some other centralized provisioning becomes a progressively smaller investment, one capable of being carried on by volunteers or by firms that need not appropriate anything approaching the full value of the product.¹³³

Moreover, integration, not only or even primarily integration into a general product but integration as a specific customization for specific users could provide an opportunity for cooperative monetary appropriation.¹³⁴ There are no models for such cooperative appropriation on a large scale yet, but the idea is that many peers will be admitted to something that is more akin to a common property regime or partnership than a commons, probably on the basis of reputation in contributing to the commons, and these groups would develop a system for receiving and disseminating service/customization projects (if it is a software project) or other information production processes. This would not necessarily work for all information production, but it could work in some. The idea is that the indirect appropriation itself would be organized on a peer model, so that reputation would lead not to being hired as an employee by a hierarchical firm, but would instead be an initiation into a cooperative, managed and “owned” by its participants. Just as in the case of Slashdot, some mechanism for assuring quality of work in the products would be necessary, but it would be achievable on a distributed, rather than a hierarchical model, with some tracking of individual contribution to any given project (or some other mechanism for distribution of revenues). The idea here would be to provide a peer-based model for allowing contributors to share the benefits of large-scale service projects, rather than relegating them to appropriation based on whatever they can individually and personally do as a result of participating in the common project.

To conclude, whether or not a peer production project will be able to resolve the integration problem is a central limiting factor on the viability of peer production to provision any given information goods. Approaches to integration include technology, as with the software running Slashdot or the clickworkers project, iterative peer production, such as the moderation and meta-moderation on Slashdot, social norms, as with Wikipedia's or Kuro5hin, and market or hierarchical mechanisms that integrate the project without appropriating the joint product, as is the case in the Linux kernel development community.

¹³³ Boyle focuses on this characteristic as the most interesting and potentially important solution. See Boyle, Second Enclosure Movement, *supra*.

¹³⁴ I owe the idea of cooperative appropriation to an enormously productive conversation with David Johnson. It was his idea that the peer production model can be combined with the producers' cooperative model to provide a mechanism of appropriation that would give contributors to peer production processes a more direct mechanism for keeping body and soul together while contributing, rather than simply awaiting for reputation gains to be translated into a contract with a company.

Conclusion

In this paper I suggest that peer production of information is a phenomenon with much broader economic implications for information production than thinking of free software alone would suggest. I describe commons-based peer production enterprises occurring throughout the value chain of information production on the Net, from content production, through relevance and accreditation, and to distribution. I then explain that peer production has a systematic advantage over markets and firms in matching the best available human capital to the best available information inputs to create information products.

Peer production of information is emerging because the declining price of physical capital involved in information production and the declining price of communications lower the cost of peer production and make human capital the primary economic good involved. This both lowers the cost of coordination and increases the importance of the factor at which peer production has a relative advantage—identifying the best available human capital in highly refined increments and allocating it to projects. If true, this would have a number of implications both for firms seeking to structure a business model for the Net, and for governments seeking to capitalize on the Net to become more innovative and productive.

For academics, peer production provides a rich area for new research. Peer production, like the Net, is just emerging. While there are some studies of peer-produced software, there is little by way of systematic research into peer production processes more generally. There is much room for theoretical work on why they work, what are potential pitfalls, and what are solutions that in principle and in practice can be adopted. The role of norms, the role of technology, and the interaction between volunteerism and economic gain in shaping the motivation and organization of peer production are also important areas of research, in particular in the study of how peer groups cluster around projects. Qualitative and quantitative studies of the importance of peer production in the overall information economy, and in particular the Internet-based information economy would provide a better picture of just how central or peripheral a phenomenon this is.

For firms, the emergence of peer production may require a more aggressive move from information product-based business models to information-embedding material products and service-based business models. Businesses could, following IBM or Red Hat in open source software, focus their “production” investment in providing opportunities for peer production, aiding in that production, and performing if necessary some of the integration functions. Firms that adopt this model, however, will not be able to count on appropriating the end product directly, because the threat

of appropriation will largely dissipate motivations for participation. Indeed, the capacity of a firm to commit credibly *not* to appropriate the joint project will be crucial to its success in building a successful business model alongside a peer production process. This would require specific licenses that secure access to the work over time to contributors and all. It would also require a business model that depends on indirect appropriation of the benefits of the product.¹³⁵ Selling products or services, for which availability of the peer-produced product increases demand, as in the case of IBM servers that run the GNU/Linux operating system and Apache server software, could do this. Conversely, firms that benefit on the supply side from access to certain types of information can capitalize on peer production processes to provide that input cheaply and efficiently, while gaining the firm-specific human capital to optimizing their product to fit the information. Again, IBM's investment in engineers who participate in writing open source software releases it from reliance on proprietary software owned by other firms, thereby creating supply side economies to its support of peer production of software. Similarly, NASA's utilization of peer production saves on its costs of mapping Mars craters, and Google's use of links provided by web sites as votes for relevance integrates distributed relevance judgments as input into its own commercial product. Another option is sale of the tools of peer production itself. For example, the software and access to massive multiplayer online games like Ultima Online or Everquest are an instance of a growing industry in the tools for peer production of escapist storytelling.¹³⁶

For regulators, the implications are quite significant. In particular, the current heavy focus on strengthening intellectual property rights is exactly the wrong approach to increasing growth through innovation and information production if having a robust peer production sector is important to an economy's capacity to tap its human capital efficiently. Strong intellectual property rights, in particular rights to control creative utilization of existing information, harm peer production by raising the cost of access to existing information resources as input. This limits the capacity of the hundreds of thousands of potential contributors to consider what could be done with a given input, and applying themselves to it without violating the rights of the owner of the information input. This does not mean that intellectual property rights are all bad. But we have known for decades that intellectual property entails systematic inefficiencies as a solution to the problem of private provisioning of the public good called information. The emergence of commons-based peer production adds a new source of inefficiency.

¹³⁵ For a general mapping of indirect appropriation mechanisms see Benkler, *Intellectual Property and the Organization of Information Production*, *supra*.

¹³⁶ See Bob Tedeschi, *E-commerce Report: The computer game industry seeks to bridge an online gap between geeks and the mainstream*, NYT (Dec. 31, 2001), C5 col. 1. A somewhat optimistic report estimates that this industry will pull in some \$1.3 billion by 2006. See *Gaming Platforms Set for Explosive Growth*, Europemedia (July 2 2002).

The strength of peer production is in matching human capital to information inputs to produce new information goods. Strong intellectual property rights inefficiently shrink the universe of existing information inputs that could be subjected to this process. Instead, owned inputs will be limited to human capital with which the owner of the input has a contractual—usually employment—relationship. Moreover, the entire universe of peer-produced information gains no benefit from strong intellectual property rights. Since the core of commons-based peer production entails provisioning without direct appropriation, and since indirect appropriation—intrinsic or extrinsic—does not rely on control of the information, but on its widest possible availability, intellectual property offers no gain, and only loss, to peer production. While it is true that free software currently uses copyright-based licensing to prevent certain kinds of defection from peer production processes, that strategy is needed only as a form of institutional Jiu-Jitsu to defend from intellectual property.¹³⁷ A complete absence of property from the software domain would be at least as congenial to free software development as the condition where property exists, but copyright permits free software projects to use licensing to defend themselves from defection. The same protection from defection might be provided by other means as well, such as creating simple public mechanisms for contributing one's work in a way that makes it unsusceptible to downstream appropriation—a conservancy of sorts. Regulators concerned with fostering innovation may better direct their efforts to provide the institutional tools that would help thousands of people to collaborate without appropriating their joint product and to make the information they produce freely available, rather than spending their efforts as they now do, to increase the scope and sophistication of the mechanisms for private appropriation of this public good.

That we cannot fully understand a phenomenon does not mean that it does not exist. That a seemingly growing phenomenon refuses to fit our settled perceptions of how people behave and how economic growth occurs counsels closer attention, not studied indifference and ignorance. Peer production presents a fascinating phenomenon that could allow us to tap substantially underutilized reserves of human creative effort. It is of central importance to policy debates today that we not squelch it, or, more likely, move its benefits to economies that do appreciate it and create the institutional conditions needed for it to flourish.

¹³⁷ But see McGowan, *supra*, at 287-88 (“Indeed, if we assume that there will always be opportunistic programmers who might try to appropriate a base of open-source code for use in a conventional program, then open-source production will always have to rely on the right to exclude being vested in a person or entity willing to wield that right to enforce community norms and thwart appropriation of the community's work.”).