The Wealth of Networks:
How Social Production Transforms Markets and Freedom
by Yochai Benkler, Yale University Press

Chapter 4
The Economics of Social Production

Epigraph

"Human nature is not a machine to be built after a model, and set to do exactly the work prescribed for it, but a tree, which requires to grow and develop itself on all sides, according to the tendency of the inward forces which make it a living thing."

"Such are the differences among human beings in their sources of pleasure, their susceptibilities of pain, and the operation on them of different physical and moral agencies, that unless there is a corresponding diversity in their modes of life, they neither obtain their fair share of happiness, nor grow up to the mental, moral, and aesthetic stature of which their nature is capable."

John Stuart Mill, On Liberty *(1859)*

Chapter 4
The Economics of Social Production

The increasing salience of nonmarket production in general, and peer production in particular, raises three puzzles from an economics perspective. First, why do people participate? What is their motivation when they work for or contribute resources to a project for which they are not paid or directly rewarded? Second, why now, why here? What, if anything, is special about the digitally networked environment that would lead us to believe that peer production is here to stay as an important economic phenomenon, as opposed to a fad that will pass as the medium matures and patterns of behavior settle toward those more familiar to us from the economy of steel, coal, and temp agencies. Third, is it efficient to have all these people sharing their computers and donating their time and creative effort? Moving through the answers to these questions, it becomes clear that the diverse and complex patterns of behavior observed on the Internet, from Viking ship hobbyists to the developers of the GNU/Linux operating system, are perfectly consistent with much of our contemporary understanding of human economic behavior. We need to assume no fundamental change in the nature of humanity; we need not declare the end of economics as we know it. We merely need to see that the material conditions of production in the networked information economy have changed in ways that increase the relative salience of social sharing and exchange as a modality of economic production. That is, behaviors and motivation patterns familiar to us from social relations generally continue to cohere in their own patterns. What has changed is that now these patterns of behavior have become effective beyond the domains of building social relations of mutual interest and fulfilling our emotional and psychological needs of companionship and mutual recognition. They have come to play a substantial role as modes of motivating, informing, and organizing productive behavior at the very core of the information economy. And it is this increasing role as a modality of information production that ripples through the rest this book. It is the feasibility of producing information, knowledge, and culture through social, rather than market and proprietary relations - through cooperative peer production and coordinate individual action - that creates the opportunities for greater autonomous action, a more critical culture, a more discursively engaged and better informed republic, and perhaps a more equitable global community.

Motivation

Much of economics achieves analytic tractability by adopting a very simple model of human motivation. The basic assumption is that all human motivations can be more or less reduced to something like positive and negative utilities - things people want, and things people want to avoid. These are capable of being summed, and are usually translatable into a universal medium of exchange, like money. Adding more of something people want, like money, to any given interaction will, all things considered, make that interaction more desirable to rational people. While simplistic, this highly tractable model of human motivation has enabled policy prescriptions that have proven far more productive than prescriptions that depended on other models of human motivation - such as assuming that benign administrators will be motivated to serve their people, or that individuals will undertake self-sacrifice for the good of the nation or the commune.

Of course, this simple model underlying much of contemporary economics is wrong. At least it is wrong as a universal description of human motivation. If you leave a fifty-dollar check on the table at the end of a dinner party at a friend's house, you do not increase the probability that you will be invited again. We live our lives in diverse social frames, and money has a complex relationship with these - sometimes it adds to the motivation to participate, sometimes it detracts from it. While this

is probably a trivial observation outside of the field of economics, it is quite radical within that analytic framework. The present generation's efforts to formalize and engage it began with the Titmuss-Arrow debate of the early 1970s. In a major work, Richard Titmuss compared the U.S. and British blood supply systems. The former was largely commercial at the time, organized by a mix of private for-profit and nonprofit actors; the latter entirely voluntary and organized by the National Health Service. Titmuss found that the British system had higher-quality blood (as measured by the likelihood of recipients contracting hepatitis from transfusions), less blood waste, and fewer blood shortages at hospitals. Titmuss also attacked the U.S. system as inequitable, arguing that the rich exploited the poor and desperate by buying their blood. He concluded that an altruistic blood procurement system is both more ethical and more efficient than a market system, and recommended that the market be kept out of blood donation to protect the "right to give."/1 Titmuss's argument came under immediate attack from economists. Most relevant for our purposes here, Kenneth Arrow agreed that the differences in blood quality indicated that the U.S. blood system was flawed, but rejected Titmuss's central theoretical claim that markets reduce donative activity. Arrow reported the alternative hypothesis held by "economists typically," that if some people respond to exhortation/moral incentives (donors), while others respond to prices and market incentives (sellers), these two groups likely behave independently - neither responds to the other's incentives. Thus, the decision to allow or ban markets should have no effect on donative behavior. Removing a market could, however, remove incentives of the "bad blood" suppliers to sell blood, thereby improving the overall quality of the blood supply. Titmuss had not established his hypothesis analytically, Arrow argued, and its proof or refutation would lie in empirical study./2 Theoretical differences aside, the U.S. blood supply system did in fact transition to an all-volunteer system of social donation since the 1970s. In surveys since, blood donors have reported that they "enjoy helping" others, experienced a sense of moral obligation or responsibility, or exhibited characteristics of reciprocators after they or their relatives received blood.

A number of scholars, primarily in psychology and economics, have attempted to resolve this question both empirically and theoretically. The most systematic work within economics is that of Swiss economist Bruno Frey and various collaborators, building on the work of psychologist Edward Deci./3 A simple statement of this model is that individuals have intrinsic and extrinsic motivations. Extrinsic motivations are imposed on individuals from the outside. They take the form of either offers of money for, or prices imposed on, behavior, or threats of punishment or reward from a manager or a judge for complying with, or failing to comply with, specifically prescribed behavior. Intrinsic motivations are reasons for action that come from within the person, such as pleasure or personal satisfaction. Extrinsic motivations are said to "crowd out" intrinsic motivations because they (a) impair self-determination - that is, people feel pressured by an external force, and therefore feel overjustified in maintaining their intrinsic motivation rather than complying with the will of the source of the extrinsic reward; or (b) impair self-esteem - they cause individuals to feel that their internal motivation is rejected, not valued, and as a result, their self-esteem is diminished, causing them to reduce effort. Intuitively, this model relies on there being a culturally contingent notion of what one "ought" to do if one is a well-adjusted human being and member of a decent society. Being offered money to do something you know you "ought" to do, and that self-respecting members of society usually in fact do, implies that the person offering the money believes that you are not a well-adjusted human being or an equally respectable member of society. This causes the person offered the money either to believe the offerer, and thereby lose self-esteem and reduce effort, or to resent him and resist the offer. A similar causal explanation is formalized by Roland Benabou and Jean Tirole, who claim that the person receiving the monetary incentives infers that the person offering the compensation does not trust the offeree to do the right thing, or to do it well of their own accord. The offeree's self-confidence and intrinsic motivation to succeed

are reduced to the extent that the offeree believes that the offerer - a manager or parent, for example - is better situated to judge the offeree's abilities./4

More powerful than the theoretical literature is the substantial empirical literature - including field and laboratory experiments, econometrics, and surveys - that has developed since the mid-1990s to test the hypotheses of this model of human motivation. Across many different settings, researchers have found substantial evidence that, under some circumstances, adding money for an activity previously undertaken without price compensation reduces, rather than increases, the level of activity. The work has covered contexts as diverse as the willingness of employees to work more or to share their experience and knowledge with team members, of communities to accept locally undesirable land uses, or of parents to pick up children from day-care centers punctually./5 The results of this empirical literature strongly suggest that across various domains some displacement or crowding out can be identified between monetary rewards and nonmonetary motivations. This does not mean that offering monetary incentives does not increase extrinsic rewards - it does. Where extrinsic rewards dominate, this will increase the activity rewarded as usually predicted in economics. However, the effect on intrinsic motivation, at least sometimes, operates in the opposite direction. Where intrinsic motivation is an important factor because pricing and contracting are difficult to achieve, or because the payment that can be offered is relatively low, the aggregate effect may be negative. Persuading experienced employees to communicate their tacit knowledge to the teams they work with is a good example of the type of behavior that is very hard to specify for efficient pricing, and therefore occurs more effectively through social motivations for teamwork than through payments. Negative effects of small payments on participation in work that was otherwise volunteer-based are an example of low payments recruiting relatively few people, but making others shift their efforts elsewhere and thereby reducing, rather than increasing, the total level of volunteering for the job.

The psychology-based alternative to the "more money for an activity will mean more of the activity" assumption implicit in most of these new economic models is complemented by a sociology-based alternative. This comes from one branch of the social capital literature - the branch that relates back to Mark Granovetter's 1974 book, *Getting a Job*, and was initiated as a crossover from sociology to economics by James Coleman./6 This line of literature rests on the claim that, as Nan Lin puts it, "there are two ultimate (or primitive) rewards for human beings in a social structure: economic standing and social standing."/7 These rewards are understood as instrumental and, in this regard, are highly amenable to economics. Both economic and social aspects represent "standing" - that is, a relational measure expressed in terms of one's capacity to mobilize resources. Some resources can be mobilized by money. Social relations can mobilize others. For a wide range of reasons - institutional, cultural, and possibly technological - some resources are more readily capable of being mobilized by social relations than by money. If you want to get your nephew a job at a law firm in the United States today, a friendly relationship with the firm's hiring partner is more likely to help than passing on an envelope full of cash. If this theory of social capital is correct, then sometimes you should be willing to trade off financial rewards for social capital. Critically, the two are not fungible or cumulative. A hiring partner paid in an economy where monetary bribes for job interviews are standard does not acquire a social obligation. That same hiring partner in that same culture, who is also a friend and therefore forgoes payment, however, probably does acquire a social obligation, tenable for a similar social situation in the future. The magnitude of the social debt, however, may now be smaller. It is likely measured by the amount of money saved from not having to pay the price, not by the value of getting the nephew a job, as it would likely be in an economy where jobs cannot be had for bribes. There are things and behaviors, then, that simply cannot be commodified for market exchange, like friendship. Any effort to mix the two, to pay for one's friendship, would render it something completely different - perhaps a psychoanalysis session in our

culture. There are things that, even if commodified, can still be used for social exchange, but the meaning of the social exchange would be diminished. One thinks of borrowing eggs from a neighbor, or lending a hand to friends who are moving their furniture to a new apartment. And there are things that, even when commodified, continue to be available for social exchange with its full force. Consider gamete donations as an example in contemporary American culture. It is important to see, though, that there is nothing intrinsic about any given "thing" or behavior that makes it fall into one or another of these categories. The categories are culturally contingent and cross-culturally diverse. What matters for our purposes here, though, is only the realization that for any given culture, there will be some acts that a person would prefer to perform not for money, but for social standing, recognition, and probably, ultimately, instrumental value obtainable only if that person has performed the action through a social, rather than a market, transaction.

It is not necessary to pin down precisely the correct or most complete theory of motivation, or the full extent and dimensions of crowding out nonmarket rewards by the introduction or use of market rewards. All that is required to outline the framework for analysis is recognition that there is some form of social and psychological motivation that is neither fungible with money nor simply cumulative with it. Transacting within the price system may either increase or decrease the social-psychological rewards (be they intrinsic or extrinsic, functional or symbolic). The intuition is simple. As I have already said, leaving a fifty-dollar check on the table after one has finished a pleasant dinner at a friend's house would not increase the host's social and psychological gains from the evening. Most likely, it would diminish them sufficiently that one would never again be invited. A bottle of wine or a bouquet of flowers would, to the contrary, improve the social gains. And if dinner is not intuitively obvious, think of sex. The point is simple. Money-oriented motivations are different from socially oriented motivations. Sometimes they align. Sometimes they collide. Which of the two will be the case is historically and culturally contingent. The presence of money in sports or entertainment reduced the social psychological gains from performance in late-nineteenth-century Victorian England, at least for members of the middle and upper classes. This is reflected in the long-standing insistence on the "amateur" status of the Olympics, or the status of "actors" in the Victorian society. This has changed dramatically more than a century later, where athletes' and popular entertainers' social standing is practically measured in the millions of dollars their performances can command.

The relative relationships of money and social-psychological rewards are, then, dependent on culture and context. Similar actions may have different meanings in different social or cultural contexts. Consider three lawyers contemplating whether to write a paper presenting their opinion - one is a practicing attorney, the second is a judge, and the third is an academic. For the first, money and honor are often, though not always, positively correlated. Being able to command a very high hourly fee for writing the requested paper is a mode of expressing one's standing in the profession, as well as a means of putting caviar on the table. Yet, there are modes of acquiring esteem - like writing the paper as a report for a bar committee - that are not improved by the presence of money, and are in fact undermined by it. This latter effect is sharpest for the judge. If a judge is approached with an offer of money for writing an opinion, not only is this not a mark of honor, it is a subversion of the social role and would render corrupt the writing of the opinion. For the judge, the intrinsic "rewards" for writing the opinion when matched by a payment for the product would be guilt and shame, and the offer therefore an expression of disrespect. Finally, if the same paper is requested of the academic, the presence of money is located somewhere in between the judge and the practitioner. To a high degree, like the judge, the academic who writes for money is rendered suspect in her community of scholarship. A paper clearly funded by a party, whose results support the party's regulatory or litigation position, is practically worthless as an academic work. In a mirror image of the practitioner, however, there are some forms of money that add to and

reinforce an academic's social psychological rewards - peer-reviewed grants and prizes most prominent among them.

Moreover, individuals are not monolithic agents. While it is possible to posit idealized avaricious money-grubbers, altruistic saints, or social climbers, the reality of most people is a composite of these all, and one that is not like any of them. Clearly, some people are more focused on making money, and others are more generous; some more driven by social standing and esteem, others by a psychological sense of well-being. The for-profit and nonprofit systems probably draw people with different tastes for these desiderata. Academic science and commercial science also probably draw scientists with similar training but different tastes for types of rewards. However, well-adjusted, healthy individuals are rarely monolithic in their requirements. We would normally think of someone who chose to ignore and betray friends and family to obtain either more money or greater social recognition as a fetishist of some form or another. We spend some of our time making money, some of our time enjoying it hedonically; some of our time being with and helping family, friends, and neighbors; some of our time creatively expressing ourselves, exploring who we are and what we would like to become. Some of us, because of economic conditions we occupy, or because of our tastes, spend very large amounts of time trying to make money - whether to become rich or, more commonly, just to make ends meet. Others spend more time volunteering, chatting, or writing.

For all of us, there comes a time on any given day, week, and month, every year and in different degrees over our lifetimes, when we choose to act in some way that is oriented toward fulfilling our social and psychological needs, not our market-exchangeable needs. It is that part of our lives and our motivational structure that social production taps, and on which it thrives. There is nothing mysterious about this. It is evident to any of us who rush home to our family or to a restaurant or bar with friends at the end of a workday, rather than staying on for another hour of overtime or to increase our billable hours; or at least regret it when we cannot. It is evident to any of us who has ever brought a cup of tea to a sick friend or relative, or received one; to anyone who has lent a hand moving a friend's belongings; played a game; told a joke, or enjoyed one told by a friend. What needs to be understood now, however, is under what conditions these many and diverse social actions can turn into an important modality of economic production. When can all these acts, distinct from our desire for money and motivated by social and psychological needs, be mobilized, directed, and made effective in ways that we recognize as economically valuable?

Social Production: Feasibility Conditions and Organizational Form

The core technologically contingent fact that enables social relations to become a salient modality of production in the networked information economy is that all the inputs necessary to effective productive activity are under the control of individual users. Human creativity, wisdom, and life experience are all possessed uniquely by individuals. The computer processors, data storage devices, and communications capacity necessary to make new meaningful conversational moves from the existing universe of information and stimuli, and to render and communicate them to others near and far are also under the control of these same individual users - at least in the advanced economies and in some portions of the population of developing economies. This does not mean that all the physical capital necessary to process, store, and communicate information is under individual user control. That is not necessary. It is, rather, that the majority of individuals in these societies have the threshold level of material capacity required to explore the information environment they occupy, to take from it, and to make their own contributions to it.

There is nothing about computation or communication that naturally or necessarily enables this fact. It is a felicitous happenstance of the fabrication technology of computing machines in the last quarter of the twentieth century, and, it seems, in the reasonably foreseeable future. It is cheaper to build freestanding computers that enable their owners to use a wide and dynamically changing range of information applications, and that are cheap enough that each machine is owned by an individual user or household, than it is to build massive supercomputers with incredibly high-speed communications to yet cheaper simple terminals, and to sell information services to individuals on an on-demand or standardized package model. Natural or contingent, it is nevertheless a fact of the industrial base of the networked information economy that individual users - susceptible as they are to acting on diverse motivations, in diverse relationships, some market-based, some social - possess and control the physical capital necessary to make effective the human capacities they uniquely and individually possess.

Now, having the core inputs of information production ubiquitously distributed in society is a core enabling fact, but it alone cannot assure that social production will become economically significant. Children and teenagers, retirees, and very rich individuals can spend most of their lives socializing or volunteering; most other people cannot. While creative capacity and judgment are universally distributed in a population, available time and attention are not, and human creative capacity cannot be fully dedicated to nonmarket, nonproprietary production all the time. Someone needs to work for money, at least some of the time, to pay the rent and put food on the table. Personal computers too are only used for earnings-generating activities some of the time. In both these resources, there remain large quantities of excess capacity - time and interest in human beings; processing, storage, and communications capacity in computers - available to be used for activities whose rewards are not monetary or monetizable, directly or indirectly.

For this excess capacity to be harnessed and become effective, the information production process must effectively integrate widely dispersed contributions, from many individual human beings and machines. These contributions are diverse in their quality, quantity, and focus, in their timing and geographic location. The great success of the Internet generally, and peer-production processes in particular, has been the adoption of technical and organizational architectures that have allowed them to pool such diverse efforts effectively. The core characteristics underlying the success of these enterprises are their modularity and their capacity to integrate many fine-grained contributions.

"Modularity" is a property of a project that describes the extent to which it can be broken down into smaller components, or modules, that can be independently 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. This maximizes their autonomy and flexibility to define the nature, extent, and timing of their participation in the project. Breaking up the maps of Mars involved in the clickworkers project (described in chapter 3) and rendering them in small segments with a simple marking tool is a way of modularizing the task of mapping craters. In the SETI@home project (see chapter 3), the task of scanning radio astronomy signals is broken down into millions of little computations as a way of modularizing the calculations involved.

"Granularity" refers to the size of the modules, in terms of the time and effort that an individual must invest in producing them. The five minutes required for moderating a comment on Slashdot, or for metamoderating a moderator, is more fine-grained than the hours necessary to participate in writing a bug fix in an open-source project. More people can participate in the former than in the latter, independent of the differences in the knowledge required for participation. The number of people who can, in principle, participate in a project is therefore inversely related to the size of the

smallest-scale contribution necessary to produce a usable module. The granularity of the modules therefore 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. Most importantly for our purposes of understanding the rising role of nonmarket production, the time can be drawn from the excess time we normally dedicate to having fun and participating in social interactions. If the finest-grained contributions are relatively large and would require a 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.

Perhaps the clearest example of how large-grained modules can make projects falter is the condition, as of the middle of 2005, of efforts to peer produce open textbooks. The largest such effort is Wikibooks, a site associated with *Wikipedia*, which has not taken off as did its famous parent project. Very few texts there have reached maturity to the extent that they could be usable as a partial textbook, and those few that have were largely written by one individual with minor contributions by others. Similarly, an ambitious initiative launched in California in 2004 still had not gone far beyond an impassioned plea for help by mid-2005. The project that seems most successful as of 2005 was a South African project, Free High School Science Texts (FHSST), founded by a physics graduate student, Mark Horner. As of this writing, that three-year-old project had more or less completed a physics text, and was about halfway through chemistry and mathematics textbooks. The whole FHSST project involves a substantially more managed approach than is common in peer-production efforts, with a core group of dedicated graduate student administrators recruiting contributors, assigning tasks, and integrating the contributions. Horner suggests that the basic limiting factor is that in order to write a high school textbook, the output must comply with state-imposed guidelines for content and form. To achieve these requirements, the various modules must cohere to a degree much larger than necessary in a project like *Wikipedia*, which can endure high diversity in style and development without losing its utility. As a result, the individual contributions have been kept at a high level of abstraction - an idea or principle explained at a time. The minimal time commitment required of each contributor is therefore large, and has led many of those who volunteered initially to not complete their contributions. In this case, the guideline requirements constrained the project's granularity, and thereby impeded its ability to grow and capture the necessary thousands of small-grained contributions. With orders of magnitude fewer contributors, each must be much more highly motivated and available than is necessary in *Wikipedia*, Slashdot, and similar successful projects.

It is not necessary, however, that each and every chunk or module be fine grained. Free software projects in particular have shown us that successful peer-production projects may also be structured, technically and culturally, in ways that make it possible for different individuals to contribute vastly different levels of effort commensurate with their ability, motivation, and availability. The large free software projects might integrate thousands of people who are acting primarily for social psychological reasons - because it is fun or cool; a few hundred young programmers aiming to make a name for themselves so as to become employable; and dozens of programmers who are paid to write free software by firms that follow one of the nonproprietary strategies described in chapter 2. IBM and Red Hat are the quintessential examples of firms that contribute paid employee time to peer-production projects in this form. This form of link between a commercial firm and a peer production community is by no means necessary for a peer-production process to succeed; it does, however, provide one constructive interface between market- and nonmarket-motivated behavior, through which actions on the two types of motivation can reinforce, rather than undermine, each other.

The characteristics of planned modularization of a problem are highly visible and explicit in some peer-production projects - the distributed computing projects like SETI@home are particularly good examples of this. However, if we were to step back and look at the entire phenomenon of Web-based publication from a bird's-eye view, we would see that the architecture of the World Wide Web, in particular the persistence of personal Web pages and blogs and their self-contained, technical independence of each other, give the Web as a whole the characteristics of modularity and variable but fine-grained granularity. Imagine that you were trying to evaluate how, if at all, the Web is performing the task of media watchdog. Consider one example, which I return to in chapter 7: The Memory Hole, a Web site created and maintained by Russ Kick, a freelance author and editor. Kick spent some number of hours preparing and filing a Freedom of Information Act request with the Defense Department, seeking photographs of coffins of U.S. military personnel killed in Iraq. He was able to do so over some period, not having to rely on "getting the scoop" to earn his dinner. At the same time, tens of thousands of other individual Web publishers and bloggers were similarly spending their time hunting down stories that moved them, or that they happened to stumble across in their own daily lives. When Kick eventually got the photographs, he could upload them onto his Web site, where they were immediately available for anyone to see. Because each contribution like Kick's can be independently created and stored, because no single permission point or failure point is present in the architecture of the Web - it is merely a way of conveniently labeling documents stored independently by many people who are connected to the Internet and use HTML (hypertext markup language) and HTTP (hypertext transfer protocol) - as an "information service," it is highly modular and diversely granular. Each independent contribution comprises as large or small an investment as its owner-operator chooses to make. Together, they form a vast almanac, trivia trove, and news and commentary facility, to name but a few, produced by millions of people at their leisure - whenever they can or want to, about whatever they want.

The independence of Web sites is what marks their major difference from more organized peer-production processes, where contributions are marked not by their independence but by their interdependence. The Web as a whole requires no formal structure of cooperation. As an "information good" or medium, it emerges as a pattern out of coordinate coexistence of millions of entirely independent acts. All it requires is a pattern recognition utility superimposed over the outputs of these acts - a search engine or directory. Peer-production processes, to the contrary, do generally require some substantive cooperation among users. A single rating of an individual comment on Slashdot does not by itself moderate the comment up or down, neither does an individual marking of a crater. Spotting a bug in free software, proposing a fix, reviewing the proposed fix, and integrating it into the software are interdependent acts that require a level of cooperation. This necessity for cooperation requires peer-production processes to adopt more engaged strategies for assuring that everyone who participates is doing so in good faith, competently, and in ways that do not undermine the whole, and weeding out those would-be participants who are not.

Cooperation in peer-production processes is usually maintained by some combination of technical architecture, social norms, legal rules, and a technically backed hierarchy that is validated by social norms. *Wikipedia* is the strongest example of a discourse-centric model of cooperation based on social norms. However, even *Wikipedia* includes, ultimately, a small number of people with system administrator privileges who can eliminate accounts or block users in the event that someone is being genuinely obstructionist. This technical fallback, however, appears only after substantial play has been given to self-policing by participants, and to informal and quasi-formal community-based dispute resolution mechanisms. Slashdot, by contrast, provides a strong model of a sophisticated technical system intended to assure that no one can "defect" from the cooperative enterprise of commenting and moderating comments. It limits behavior enabled by the system to avoid

destructive behavior before it happens, rather than policing it after the fact. The Slash code does this by technically limiting the power any given person has to moderate anyone else up or down, and by making every moderator the subject of a peer review system whose judgments are enforced technically - that is, when any given user is described by a sufficiently large number of other users as unfair, that user automatically loses the technical ability to moderate the comments of others. The system itself is a free software project, licensed under the GPL (General Public License) - which is itself the quintessential example of how law is used to prevent some types of defection from the common enterprise of peer production of software. The particular type of defection that the GPL protects against is appropriation of the joint product by any single individual or firm, the risk of which would make it less attractive for anyone to contribute to the project to begin with. The GPL assures that, as a legal matter, no one who contributes to a free software project need worry that some other contributor will take the project and make it exclusively their own. The ultimate quality judgments regarding what is incorporated into the "formal" releases of free software projects provide the clearest example of the extent to which a meritocratic hierarchy can be used to integrate diverse contributions into a finished single product. In the case of the Linux kernel development project (see chapter 3), it was always within the power of Linus Torvalds, who initiated the project, to decide which contributions should be included in a new release, and which should not. But it is a funny sort of hierarchy, whose quirkiness Steve Weber well explicates./8 Torvalds's authority is persuasive, not legal or technical, and certainly not determinative. He can do nothing except persuade others to prevent them from developing anything they want and add it to their kernel, or to distribute that alternative version of the kernel. There is nothing he can do to prevent the entire community of users, or some subsection of it, from rejecting his judgment about what ought to be included in the kernel. Anyone is legally free to do as they please. So these projects are based on a hierarchy of meritocratic respect, on social norms, and, to a great extent, on the mutual recognition by most players in this game that it is to everybody's advantage to have someone overlay a peer review system with some leadership.

In combination then, three characteristics make possible the emergence of information production that is not based on exclusive proprietary claims, not aimed toward sales in a market for either motivation or information, and not organized around property and contract claims to form firms or market exchanges. First, the physical machinery necessary to participate in information and cultural production is almost universally distributed in the population of the advanced economies. Certainly, personal computers as capital goods are under the control of numbers of individuals that are orders of magnitude larger than the number of parties controlling the use of mass-production-capable printing presses, broadcast transmitters, satellites, or cable systems, record manufacturing and distribution chains, and film studios and distribution systems. This means that the physical machinery can be put in service and deployed in response to any one of the diverse motivations individual human beings experience. They need not be deployed in order to maximize returns on the financial capital, because financial capital need not be mobilized to acquire and put in service any of the large capital goods typical of the industrial information economy. Second, the primary raw materials in the information economy, unlike the industrial economy, are public goods - existing information, knowledge, and culture. Their actual marginal social cost is zero. Unless regulatory policy makes them purposefully expensive in order to sustain the proprietary business models, acquiring raw materials also requires no financial capital outlay. Again, this means that these raw materials can be deployed for any human motivation. They need not maximize financial returns. Third, the technical architectures, organizational models, and social dynamics of information production and exchange on the Internet have developed so that they allow us to structure the solution to problems - in particular to information production problems - in ways that are highly modular. This allows many diversely motivated people to act for a wide range of reasons that, in combination, cohere into new useful information, knowledge, and cultural goods. These

architectures and organizational models allow both independent creation that coexists and coheres into usable patterns, and interdependent cooperative enterprises in the form of peer-production processes.

Together, these three characteristics suggest that the patterns of social production of information that we are observing in the digitally networked environment are not a fad. They are, rather, a sustainable pattern of human production given the characteristics of the networked information economy. The diversity of human motivation is nothing new. We now have a substantial literature documenting its importance in free and open-source software development projects, from Josh Lerner and Jean Tirole, Rishab Ghosh, Eric Von Hippel and Karim Lakhani, and others. Neither is the public goods nature of information new. What is new are the technological conditions that allow these facts to provide the ingredients of a much larger role in the networked information economy for nonmarket, nonproprietary production to emerge. As long as capitalization and ownership of the physical capital base of this economy remain widely distributed and as long as regulatory policy does not make information inputs artificially expensive, individuals will be able to deploy their own creativity, wisdom, conversational capacities, and connected computers, both independently and in loose interdependent cooperation with others, to create a substantial portion of the information environment we occupy. Moreover, we will be able to do so for whatever reason we choose - through markets or firms to feed and clothe ourselves, or through social relations and open communication with others, to give our lives meaning and context.

Transaction Costs and Efficiency

For purposes of analyzing the political values that are the concern of most of this book, all that is necessary is that we accept that peer production in particular, and nonmarket information production and exchange in general, are sustainable in the networked information economy. Most of the remainder of the book seeks to evaluate why, and to what extent, the presence of a substantial nonmarket, commons-based sector in the information production system is desirable from the perspective of various aspects of freedom and justice. Whether this sector is "efficient" within the meaning of the word in welfare economics is beside the point to most of these considerations. Even a strong commitment to a pragmatic political theory, one that accepts and incorporates into its consideration the limits imposed by material and economic reality, need not aim for "efficient" policy in the welfare sense. It is sufficient that the policy is economically and socially sustainable on its own bottom - in other words, that it does not require constant subsidization at the expense of some other area excluded from the analysis. It is nonetheless worthwhile spending a few pages explaining why, and under what conditions, commons-based peer production, and social production more generally, are not only sustainable but actually efficient ways of organizing information production.

The efficient allocation of two scarce resources and one public good are at stake in the choice between social production - whether it is peer production or independent nonmarket production- and market-based production. Because most of the outputs of these processes are nonrival goods - information, knowledge, and culture - the fact that the social production system releases them freely, without extracting a price for using them, means that it would, all other things being equal, be more efficient for information to be produced on a nonproprietary social model, rather than on a proprietary market model. Indeed, all other things need not even be equal for this to hold. It is enough that the net value of the information produced by commons-based social production processes and released freely for anyone to use as they please is no less than the total value of information produced through property-based systems minus the deadweight loss caused by the above-marginal-cost pricing practices that are the intended result of the intellectual property system.

The two scarce resources are: first, human creativity, time, and attention; and second, the computation and communications resources used in information production and exchange. In both cases, the primary reason to choose among proprietary and nonproprietary strategies, between market-based systems - be they direct market exchange or firm-based hierarchical production - and social systems, are the comparative transaction costs of each, and the extent to which these transaction costs either outweigh the benefits of working through each system, or cause the system to distort the information it generates so as to systematically misallocate resources.

The first thing to recognize is that markets, firms, and social relations are three distinct transactional frameworks. Imagine that I am sitting in a room and need paper for my printer. I could (a) order paper from a store; (b) call the storeroom, if I am in a firm or organization that has one, and ask the clerk to deliver the paper I need; or (c) walk over to a neighbor and borrow some paper. Choice (a) describes the market transactional framework. The store knows I need paper immediately because I am willing to pay for it now. Alternative (b) is an example of the firm as a transactional framework. The paper is in the storeroom because someone in the organization planned that someone else would need paper today, with some probability, and ordered enough to fill that expected need. The clerk in the storeroom gives it to me because that is his job; again, defined by someone who planned to have someone available to deliver paper when someone else in the proper channels of authority says that she needs it. Comparing and improving the efficiency of (a) and (b), respectively, has been a central project in transaction-costs organization theory. We might compare, for example, the costs of taking my call, verifying the credit card information, and sending a delivery truck for my one batch of paper, to the costs of someone planning for the average needs of a group of people like me, who occasionally run out of paper, and stocking a storeroom with enough paper and a clerk to fill our needs in a timely manner. However, notice that (c) is also an alternative transactional framework. I could, rather than incurring the costs of transacting through the market with the local store or of building a firm with sufficient lines of authority to stock and manage the storeroom, pop over to my neighbor and ask for some paper. This would make sense even within an existing firm when, for example, I need two or three pages immediately and do not want to wait for the storeroom clerk to do his rounds, or more generally, if I am working at home and the costs of creating "a firm," stocking a storeroom, and paying a clerk are too high for my neighbors and me. Instead, we develop a set of neighborly social relations, rather than a firm-based organization, to deal with shortfalls during periods when it would be too costly to assure a steady flow of paper from the market - for example, late in the evening, on a weekend, or in a sparsely populated area.

The point is not, of course, to reduce all social relations and human decency to a transaction-costs theory. Too many such straight planks have already been cut from the crooked timber of humanity to make that exercise useful or enlightening. The point is that most of economics internally has been ignoring the social transactional framework as an alternative whose relative efficiency can be accounted for and considered in much the same way as the relative cost advantages of simple markets when compared to the hierarchical organizations that typify much of our economic activity-firms.

A market transaction, in order to be efficient, must be clearly demarcated as to what it includes, so that it can be priced efficiently. That price must then be paid in equally crisply delineated currency. Even if a transaction initially may be declared to involve sale of "an amount reasonably required to produce the required output," for a "customary" price, at some point what was provided and what is owed must be crystallized and fixed for a formal exchange. The crispness is a functional requirement of the price system. It derives from the precision and formality of the medium of exchange - currency - and the ambition to provide refined representations of the comparative value of marginal decisions through denomination in an exchange medium that represents these

incremental value differences. Similarly, managerial hierarchies require a crisp definition of who should be doing what, when, and how, in order to permit the planning and coordination process to be effective.

Social exchange, on the other hand, does not require the same degree of crispness at the margin. As Maurice Godelier put it in *The Enigma of the Gift*, "the mark of the gift between close friends and relatives . . . is not the absence of obligations, it is the absence of 'calculation.' "/9 There are, obviously, elaborate and formally ritualistic systems of social exchange, in both ancient societies and modern. There are common-property regimes that monitor and record calls on the common pool very crisply. However, in many of the common-property regimes, one finds mechanisms of bounding or fairly allocating access to the common pool that more coarsely delineate the entitlements, behaviors, and consequences than is necessary for a proprietary system. In modern market society, where we have money as a formal medium of precise exchange, and where social relations are more fluid than in traditional societies, social exchange certainly occurs as a fuzzier medium. Across many cultures, generosity is understood as imposing a debt of obligation; but none of the precise amount of value given, the precise nature of the debt to be repaid, or the date of repayment need necessarily be specified. Actions enter into a cloud of goodwill or membership, out of which each agent can understand him- or herself as being entitled to a certain flow of dependencies or benefits in exchange for continued cooperative behavior. This may be an ongoing relationship between two people, a small group like a family or group of friends, and up to a general level of generosity among strangers that makes for a decent society. The point is that social exchange does not require defining, for example, "I will lend you my car and help you move these five boxes on Monday, and in exchange you will feed my fish next July," in the same way that the following would: "I will move five boxes on Tuesday for $100, six boxes for $120." This does not mean that social systems are cost free - far from it. They require tremendous investment, acculturation, and maintenance. This is true in this case every bit as much as it is true for markets or states. Once functional, however, social exchanges require less information crispness at the margin.

Both social and market exchange systems require large fixed costs - the setting up of legal institutions and enforcement systems for markets, and creating social networks, norms, and institutions for the social exchange. Once these initial costs have been invested, however, market transactions systematically require a greater degree of precise information about the content of actions, goods, and obligations, and more precision of monitoring and enforcement on a per-transaction basis than do social exchange systems.

This difference between markets and hierarchical organizations, on the one hand, and peer-production processes based on social relations, on the other, is particularly acute in the context of human creative labor - one of the central scarce resources that these systems must allocate in the networked information economy. The levels and focus of individual effort are notoriously hard to specify for pricing or managerial commands, considering all aspects of individual effort and ability - talent, motivation, workload, and focus - as they change in small increments over the span of an individual's full day, let alone months. What we see instead is codification of effort types - a garbage collector, a law professor - that are priced more or less finely. However, we only need to 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*. Pricing therefore continues to be a function of relatively crude information about the actual variability among people. More importantly, as aspects of performance that are harder to fully specify in advance or monitor-like creativity over time given the occurrence of new

opportunities to be creative, or implicit know-how - become a more significant aspect of what is valuable about an individual's contribution, market mechanisms become more and more costly to maintain efficiently, and, as a practical matter, simply lose a lot of information.

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 the weight of human intellectual effort increases in the overall mix of inputs into a given production process, an organization model that does not require contractual specification of the individual effort required to participate in a collective enterprise, and which 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. Some firms try to solve this problem by utilizing market- and social-relations-oriented hybrids, like incentive compensation schemes and employee-of-the-month-type social motivational frameworks. These may be able to improve on firm-only or market-only approaches. It is unclear, though, how well they can overcome the core difficulty: that is, that both markets and firm hierarchies 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 in a commons-based effort. It is that the widely distributed model of information production will better identify the best person to produce a specific component of a project, considering all abilities and availability to work on the specific module within a specific time frame. 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, freedom of action for individuals coupled with 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 actions at a given time. Markets and firm incentive schemes are aimed at producing precisely this form of self-identification. However, the rigidities associated with collecting and comprehending bids from individuals through these systems (that is, transaction costs) limit the efficacy of self-identification by comparison to a system in which, once an individual self-identifies for a task, he or she can then undertake it without permission, contract, or instruction from another. The emergence of networked organizations (described and analyzed in the work of Charles Sabel and others) suggests that firms are in fact trying to overcome these limitations by developing parallels to the freedom to learn, innovate, and act on these innovations that is intrinsic to peer-production processes by loosening the managerial bonds, locating more of the conception and execution of problem solving away from the managerial core of the firm, and implementing these through social, as well as monetary, motivations. However, the need to assure that the value created is captured within the organization limits the extent to which these strategies can be implemented within a single enterprise, as opposed to their implementation in an open process of social production. This effect, in turn, is in some sectors attenuated through the use of what Walter Powell and others have described as learning networks. Engineers and scientists often create frameworks that allow them to step out of their organizational affiliations, through conferences or workshops. By reproducing the social production characteristics of academic exchange, they overcome some of the information loss caused by the boundary of the firm. While these organizational strategies attenuate the problem, they also underscore the degree to which it is widespread and understood by organizations as such. The fact that the direction of the solutions business organizations choose tends to shift elements of the production process away from market- or firm-based models and toward networked social production models is revealing. Now, the self-identification that is central to the relative

information efficiency of peer production is not always perfect. 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 *Wikipedia* or Slashdot, or as redundancy and statistical averaging do in the case of NASA clickworkers. The prevalence of misperceptions that individual contributors 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. They parallel quality control problems faced by firms and markets.

The lack of crisp specification of who is giving what to whom, and in exchange for what, also bears on the comparative transaction costs associated with the allocation of the second major type of scarce resource in the networked information economy: the physical resources that make up the networked information environment - communications, computation, and storage capacity. It is important to note, however, that these are very different from creativity and information as inputs: they are private goods, not a public good like information, and they are standardized goods with well-specified capacities, not heterogeneous and highly uncertain attributes like human creativity at a given moment and context. Their outputs, unlike information, are not public goods. The reasons that they are nonetheless subject to efficient sharing in the networked environment therefore require a different economic explanation. However, the sharing of these material resources, like the sharing of human creativity, insight, and attention, nonetheless relies on both the comparative transaction costs of markets and social relations and the diversity of human motivation.

Personal computers, wireless transceivers, and Internet connections are "shareable goods." The basic intuition behind the concept of shareable goods is simple. There are goods that are "lumpy": given a state of technology, they can only be produced in certain discrete bundles that offer discontinuous amounts of functionality or capacity. In order to have any ability to run a computation, for example, a consumer must buy a computer processor. These, in turn, only come in discrete units with a certain speed or capacity. One could easily imagine a world where computers are very large and their owners sell computation capacity to consumers "on demand," whenever they needed to run an application. That is basically the way the mainframe world of the 1960s and 1970s worked. However, the economics of microchip fabrication and of network connections over the past thirty years, followed by storage technology, have changed that. For most functions that users need, the price-performance trade-off favors stand-alone, general-purpose personal computers, owned by individuals and capable of running locally most applications users want, over remote facilities capable of selling on-demand computation and storage. So computation and storage today come in discrete, lumpy units. You can decide to buy a faster or slower chip, or a larger or smaller hard drive, but once you buy them, you have the capacity of these machines at your disposal, whether you need it or not.

Lumpy goods can, in turn, be fine-, medium-, or large-grained. A large-grained good is one that is so expensive it can only be used by aggregating demand for it. Industrial capital equipment, like a steam engine, is of this type. Fine-grained goods are of a granularity that allows consumers to buy precisely as much of the goods needed for the amount of capacity they require. Medium-grained goods are small enough for an individual to justify buying for her own use, given their price and her willingness and ability to pay for the functionality she plans to use. A personal computer is a medium-grained lumpy good in the advanced economies and among the more well-to-do in poorer countries, but is a large-grained capital good for most people in poor countries. If, given the price of such a good and the wealth of a society, a large number of individuals buy and use such medium-grained lumpy goods, that society will have a large amount of excess capacity "out there,"

in the hands of individuals. Because these machines are put into service to serve the needs of individuals, their excess capacity is available for these individuals to use as they wish - for their own uses, to sell to others, or to share with others. It is the combination of the fact that these machines are available at prices (relative to wealth) that allow users to put them in service based purely on their value for personal use, and the fact that they have enough capacity to facilitate additionally the action and fulfill the needs of others, that makes them "shareable." If they were so expensive that they could only be bought by pooling the value of a number of users, they would be placed in service either using some market mechanism to aggregate that demand, or through formal arrangements of common ownership by all those whose demand was combined to invest in purchasing the resource. If they were so finely grained in their capacity that there would be nothing left to share, again, sharing would be harder to sustain. The fact that they are both relatively inexpensive and have excess capacity makes them the basis for a stable model of individual ownership of resources combined with social sharing of that excess capacity.

Because social sharing requires less precise specification of the transactional details with each transaction, it has a distinct advantage over market-based mechanisms for reallocating the excess capacity of shareable goods, particularly when they have small quanta of excess capacity relative to the amount necessary to achieve the desired outcome. For example, imagine that there are one thousand people in a population of computer owners. Imagine that each computer is capable of performing one hundred computations per second, and that each computer owner needs to perform about eighty operations per second. Every owner, in other words, has twenty operations of excess capacity every second. Now imagine that the marginal transaction costs of arranging a sale of these twenty operations - exchanging PayPal (a widely used low-cost Internet-based payment system) account information, insurance against nonpayment, specific statement of how much time the computer can be used, and so forth - cost ten cents more than the marginal transaction costs of sharing the excess capacity socially. John wants to render a photograph in one second, which takes two hundred operations per second. Robert wants to model the folding of proteins, which takes ten thousand operations per second. For John, a sharing system would save fifty cents - assuming he can use his own computer for half of the two hundred operations he needs. He needs to transact with five other users to "rent" their excess capacity of twenty operations each. Robert, on the other hand, needs to transact with five hundred individual owners in order to use their excess capacity, and for him, using a sharing system is fifty dollars cheaper. The point of the illustration is simple. The cost advantage of sharing as a transactional framework relative to the price system increases linearly with the number of transactions necessary to acquire the level of resources necessary for an operation. If excess capacity in a society is very widely distributed in small dollops, and for any given use of the excess capacity it is necessary to pool the excess capacity of thousands or even millions of individual users, the transaction-cost advantages of the sharing system become significant.

The transaction-cost effect is reinforced by the motivation crowding out theory. When many discrete chunks of excess capacity need to be pooled, each distinct contributor cannot be paid a very large amount. Motivation crowding out theory would predict that when the monetary rewards to an activity are low, the negative effect of crowding out the social-psychological motivation will weigh more heavily than any increased incentive that is created by the promise of a small payment to transfer one's excess capacity. The upshot is that when the technological state results in excess capacity of physical capital being widely distributed in small dollops, social sharing can outperform secondary markets as a mechanism for harnessing that excess capacity. This is so because of both transaction costs and motivation. Fewer owners will be willing to sell their excess capacity cheaply than to give it away for free in the right social context and the transaction costs of selling will be higher than those of sharing.

From an efficiency perspective, then, there are clear reasons to think that social production systems - both peer production of information, knowledge, and culture and sharing of material resources - can be more efficient than market-based systems to motivate and allocate both human creative effort and the excess computation, storage, and communications capacity that typify the networked information economy. That does not mean that all of us will move out of market-based productive relationships all of the time. It does mean that alongside our market-based behaviors we generate substantial amounts of human creativity and mechanical capacity. The transaction costs of clearing those resources through the price system or through firms are substantial, and considerably larger for the marginal transaction than clearing them through social-sharing mechanisms as a transactional framework. With the right institutional framework and peer-review or quality-control mechanisms, and with well-modularized organization of work, social sharing is likely to identify the best person available for a job and make it feasible for that person to work on that job using freely available information inputs. Similarly, social transactional frameworks are likely to be substantially less expensive than market transactions for pooling large numbers of discrete, small increments of the excess capacity of the personal computer processors, hard drives, and network connections that make up the physical capital base of the networked information economy. In both cases, given that much of what is shared is excess capacity from the perspective of the contributors, available to them after they have fulfilled some threshold level of their market-based consumption requirements, social-sharing systems are likely to tap in to social psychological motivations that money cannot tap, and, indeed, that the presence of money in a transactional framework could nullify. Because of these effects, social sharing and collaboration can provide not only a sustainable alternative to market-based and firm-based models of provisioning information, knowledge, culture, and communications, but also an alternative that more efficiently utilizes the human and physical capital base of the networked information economy. A society whose institutional ecology permitted social production to thrive would be more productive under these conditions than a society that optimized its institutional environment solely for market- and firm-based production, ignoring its detrimental effects to social production.

The Emergence of Social Production in the Digitally Networked Environment

There is a curious congruence between the anthropologists of the gift and mainstream economists today. Both treat the gift literature as being about the periphery, about societies starkly different from modern capitalist societies. As Godelier puts it, "What a contrast between these types of society, these social and mental universes, and today's capitalist society where the majority of social relations are impersonal (involving the individual as citizen and the state, for instance), and where the exchange of things and services is conducted for the most part in an anonymous marketplace, leaving little room for an economy and moral code based on gift-giving."/10 And yet, sharing is everywhere around us in the advanced economies. Since the 1980s, we have seen an increasing focus, in a number of literatures, on production practices that rely heavily on social rather than price-based or governmental policies. These include, initially, the literature on social norms and social capital, or trust./11 Both these lines of literature, however, are statements of the institutional role of social mechanisms for enabling market exchange and production. More direct observations of social production and exchange systems are provided by the literature on social provisioning of public goods - like social norm enforcement as a dimension of policing criminality, and the literature on common property regimes./12 The former are limited by their focus on public goods provisioning. The latter are usually limited by their focus on discretely identifiable types of resources-common pool resources - that must be managed as among a group of claimants while retaining a proprietary outer boundary toward nonmembers. The focus of those who study these phenomena is usually on relatively small and tightly knit communities, with clear boundaries between members and nonmembers./13

These lines of literature point to an emerging understanding of social production and exchange as an alternative to markets and firms. Social production is not limited to public goods, to exotic, out-of-the-way places like surviving medieval Spanish irrigation regions or the shores of Maine's lobster fishing grounds, or even to the ubiquitous phenomenon of the household. As SETI@home and Slashdot suggest, it is not necessarily limited to stable communities of individuals who interact often and know each other, or who expect to continue to interact personally. Social production of goods and services, both public and private, is ubiquitous, though unnoticed. It sometimes substitutes for, and sometimes complements, market and state production everywhere. It is, to be fanciful, the dark matter of our economic production universe.

Consider the way in which the following sentences are intuitively familiar, yet as a practical matter, describe the provisioning of goods or services that have well-defined NAICS categories (the categories used by the Economic Census to categorize economic sectors) whose provisioning through the markets is accounted for in the Economic Census, but that are commonly provisioned in a form consistent with the definition of sharing - on a radically distributed model, without price or command.

NAICS 624410624410 [Babysitting services, child day care]

> "John, could you pick up Bobby today when you take Lauren to soccer? I have a conference call I have to make."
>
> "Are you doing homework with Zoe today, or shall I?"

NAICS 484210 [Trucking used household, office, or institutional furniture and equipment]

> "Jane, could you lend a hand moving this table to the dining room?"
>
> "Here, let me hold the elevator door for you, this looks heavy."

NAICS 484122 [Trucking, general freight, long-distance, less-than-truckload]

> "Jack, do you mind if I load my box of books in your trunk so you can drop it off at my brother's on your way to Boston?"

NAICS 514110 [Traffic reporting services]

> "Oh, don't take I-95, it's got horrible construction traffic to exit 39."

NAICS 711510 [Newspaper columnists, independent (freelance)]

> "I don't know about Kerry, he doesn't move me, I think he should be more aggressive in criticizing Bush on Iraq."

NAICS 621610 [Home health-care services]

> "Can you please get me my medicine? I'm too wiped to get up."
>
> "Would you like a cup of tea?"

NAICS 561591 [Tourist information bureaus]

> "Excuse me, how do I get to Carnegie Hall?"

NAICS 561321 [Temporary help services]

> "I've got a real crunch on the farm, can you come over on Saturday and lend a hand?"

"This is crazy, I've got to get this document out tonight, could you lend me a hand with proofing and pulling it all together tonight?"

NAICS 71 [Arts, entertainment, and recreation]

"Did you hear the one about the Buddhist monk, the Rabbi, and the Catholic priest . . . ?"

"Roger, bring out your guitar. . . . "

"Anybody up for a game of . . . ?"

The litany of examples generalizes through a combination of four dimensions that require an expansion from the current focus of the literatures related to social production. First, they relate to production of goods and services, not only of norms or rules. Social relations provide the very motivations for, and information relating to, production and exchange, not only the institutional framework for organizing action, which itself is motivated, informed, and organized by markets or managerial commands. Second, they relate to all kinds of goods, not only public goods. In particular, the paradigm cases of free software development and distributed computing involve labor and shareable goods - each plainly utilizing private goods as inputs, and, in the case of distributed computing, producing private goods as outputs. Third, at least some of them relate not only to relations of production within well-defined communities of individuals who have repeated interactions, but extend to cover baseline standards of human decency. These enable strangers to ask one another for the time or for directions, enable drivers to cede the road to each other, and enable strangers to collaborate on software projects, on coauthoring an online encyclopedia, or on running simulations of how proteins fold. Fourth, they may either complement or substitute for market and state production systems, depending on the social construction of mixed provisioning. It is hard to measure the weight that social and sharing-based production has in the economy. Our intuitions about capillary systems would suggest that the total volume of boxes or books moved or lifted, instructions given, news relayed, and meals prepared by family, friends, neighbors, and minimally decent strangers would be very high relative to the amount of substitutable activity carried on through market exchanges or state provisioning.

Why do we, despite the ubiquity of social production, generally ignore it as an economic phenomenon, and why might we now reconsider its importance? A threshold requirement for social sharing to be a modality of economic production, as opposed to one purely of social reproduction, is that sharing-based action be effective. Efficacy of individual action depends on the physical capital requirements for action to become materially effective, which, in turn, depend on technology. Effective action may have very low physical capital requirements, so that every individual has, by natural capacity, "the physical capital" necessary for action. Social production or sharing can then be ubiquitous (though in practice, it may not). Vocal cords to participate in a sing-along or muscles to lift a box are obvious examples. When the capital requirements are nontrivial, but the capital good is widely distributed and available, sharing can similarly be ubiquitous and effective. This is true both when the shared resource or good is the capacity of the capital good itself - as in the case of shareable goods - and when some widely distributed human capacity is made effective through the use of the widely distributed capital goods - as in the case of human creativity, judgment, experience, and labor shared in online peer-production processes - in which participants contribute using the widespread availability of connected computers. When use of larger-scale physical capital goods is a threshold requirement of effective action, we should not expect to see widespread reliance on decentralized sharing as a standard modality of production. Industrial mass-manufacture of automobiles, steel, or plastic toys, for example, is not the sort of thing that is likely to be produced on a social-sharing basis, because of the capital constraints. This is not to say that even for large-scale capital projects, like irrigation systems and dams, social production systems

cannot step into the breach. We have those core examples in the common-property regime literature, and we have worker-owned firms as examples of mixed systems. However, those systems tend to replicate the characteristics of firm, state, or market production - using various combinations of quotas, scrip systems, formal policing by "professional" officers, or management within worker-owned firms. By comparison, the "common property" arrangements described among lobster gangs of Maine or fishing groups in Japan, where capital requirements are much lower, tend to be more social-relations-based systems, with less formalized or crisp measurement of contributions to, and calls on, the production system.

To say that sharing is technology dependent is not to deny that it is a ubiquitous human phenomenon. Sharing is so deeply engrained in so many of our cultures that it would be difficult to argue that with the "right" (or perhaps "wrong") technological contingencies, it would simply disappear. My claim, however, is narrower. It is that the relative economic role of sharing changes with technology. There are technological conditions that require more or less capital, in larger or smaller packets, for effective provisioning of goods, services, and resources the people value. As these conditions change, the relative scope for social-sharing practices to play a role in production changes. When goods, services, and resources are widely dispersed, their owners can choose to engage with each other through social sharing instead of through markets or a formal, state-based relationship, because individuals have available to them the resources necessary to engage in such behavior without recourse to capital markets or the taxation power of the state. If technological changes make the resources necessary for effective action rare or expensive, individuals may wish to interact in social relations, but they can now only do so ineffectively, or in different fields of endeavor that do not similarly require high capitalization. Large-packet, expensive physical capital draws the behavior into one or the other of the modalities of production that can collect the necessary financial capital - through markets or taxation. Nothing, however, prevents change from happening in the opposite direction. Goods, services, and resources that, in the industrial stage of the information economy required large-scale, concentrated capital investment to provision, are now subject to a changing technological environment that can make sharing a better way of achieving the same results than can states, markets, or their hybrid, regulated industries.

Because of changes in the technology of the industrial base of the most advanced economies, social sharing and exchange is becoming a common modality of production at their very core - in the information, culture, education, computation, and communications sectors. Free software, distributed computing, ad hoc mesh wireless networks, and other forms of peer production offer clear examples of large-scale, measurably effective sharing practices. The highly distributed capital structure of contemporary communications and computation systems is largely responsible for this increased salience of social sharing as a modality of economic production in that environment. By lowering the capital costs required for effective individual action, these technologies have allowed various provisioning problems to be structured in forms amenable to decentralized production based on social relations, rather than through markets or hierarchies.

My claim is not, of course, that we live in a unique moment of humanistic sharing. It is, rather, that our own moment in history suggests a more general observation. The technological state of a society, in particular the extent to which individual agents can engage in efficacious production activities with material resources under their individual control, affects the opportunities for, and hence the comparative prevalence and salience of, social, market - both price-based and managerial - and state production modalities. The capital cost of effective economic action in the industrial economy shunted sharing to its economic peripheries - to households in the advanced economies, and to the global economic peripheries that have been the subject of the anthropology of gift or the common-property regime literatures. The emerging restructuring of capital investment in digital

networks - in particular, the phenomenon of user-capitalized computation and communications capabilities - are at least partly reversing that effect. Technology does not determine the level of sharing. It does, however, set threshold constraints on the effective domain of sharing as a modality of economic production. Within the domain of the practically feasible, the actual level of sharing practices will be culturally driven and cross-culturally diverse.

Most practices of production - social or market-based - are already embedded in a given technological context. They present no visible "problem" to solve or policy choice to make. We do not need to be focused consciously on improving the conditions under which friends lend a hand to each other to move boxes, make dinner, or take kids to school. We feel no need to reconsider the appropriateness of market-based firms as the primary modality for the production of automobiles. However, in moments where a field of action is undergoing a technological transition that changes the opportunities for sharing as a modality of production, understanding that sharing *is* a modality of production becomes more important, as does understanding how it functions as such. This is so, as we are seeing today, when prior technologies have already set up market- or state-based production systems that have the law and policy-making systems already designed to fit their requirements. While the prior arrangement may have been the most efficient, or even may have been absolutely necessary for the incumbent production system, its extension under new technological conditions may undermine, rather than improve, the capacity of a society to produce and provision the goods, resources, or capacities that are the object of policy analysis. This is, as I discuss in part III, true of wireless communications regulation, or "spectrum management," as it is usually called; of the regulation of information, knowledge, and cultural production, or "intellectual property," as it is usually now called; and it may be true of policies for computation and wired communications networks, as distributed computing and the emerging peer-to-peer architectures suggest.

The Interface of Social Production and Market-Based Businesses

The rise of social production does not entail a decline in market-based production. Social production first and foremost harnesses impulses, time, and resources that, in the industrial information economy, would have been wasted or used purely for consumption. Its immediate effect is therefore likely to increase overall productivity in the sectors where it is effective. But that does not mean that its effect on market-based enterprises is neutral. A newly effective form of social behavior, coupled with a cultural shift in tastes as well as the development of new technological and social solution spaces to problems that were once solved through market-based firms, exercises a significant force on the shape and conditions of market action. Understanding the threats that these developments pose to some incumbents explains much of the political economy of law in this area, which will occupy chapter 11. At the simplest level, social production in general and peer production in particular present new sources of competition to incumbents that produce information goods for which there are now socially produced substitutes. Open source software development, for example, first received mainstream media attention in 1998 due to publication of a leaked internal memorandum from Microsoft, which came to be known as The Halloween Memo. In it, a Microsoft strategist identified the open source methodology as the one major potential threat to the company's dominance over the desktop. As we have seen since, definitively in the Web server market and gradually in segments of the operating system market, this prediction proved prescient. Similarly, *Wikipedia* now presents a source of competition to online encyclopedias like *Columbia*, *Grolier*, or *Encarta*, and may well come to be seen as an adequate substitute for *Britannica* as well. Most publicly visible, peer-to-peer file sharing networks have come to compete with the recording industry as an alternative music distribution system, to the point where the long-term existence of that industry is in question. Some scholars like William Fisher, and artists like

Jenny Toomey and participants in the Future of Music Coalition, are already looking for alternative ways of securing for artists a living from the music they make.

The competitive threat from social production, however, is merely a surface phenomenon. Businesses often face competition or its potential, and this is a new source, with new economics, which may or may not put some of the incumbents out of business. But there is nothing new about entrants with new business models putting slow incumbents out of business. More basic is the change in opportunity spaces, the relationships of firms to users, and, indeed, the very nature of the boundary of the firm that those businesses that are already adapting to the presence and predicted persistence of social production are exhibiting. Understanding the opportunities social production presents for businesses begins to outline how a stable social production system can coexist and develop a mutually reinforcing relationship with market-based organizations that adapt to and adopt, instead of fight, them.

Consider the example I presented in chapter 2 of IBM's relationship to the free and open source software development community. IBM, as I explained there, has shown more than $2 billion a year in "Linux-related revenues." Prior to IBM's commitment to adapting to what the firm sees as the inevitability of free and open source software, the company either developed in house or bought from external vendors the software it needed as part of its hardware business, on the one hand, and its software services - customization, enterprise solutions, and so forth - on the other hand. In each case, the software development follows a well-recognized supply chain model. Through either an employment contract or a supply contract the company secures a legal right to require either an employee or a vendor to deliver a given output at a given time. In reliance on that notion of a supply chain that is fixed or determined by a contract, the company turns around and promises to its clients that it will deliver the integrated product or service that includes the contracted-for component. With free or open source software, that relationship changes. IBM is effectively relying for its inputs on a loosely defined cloud of people who are engaged in productive social relations. It is making the judgment that the probability that a sufficiently good product will emerge out of this cloud is high enough that it can undertake a contractual obligation to its clients, even though no one in the cloud is specifically contractually committed to it to produce the specific inputs the firm needs in the timeframe it needs it. This apparent shift from a contractually deterministic supply chain to a probabilistic supply chain is less dramatic, however, than it seems. Even when contracts are signed with employees or suppliers, they merely provide a probability that the employee or the supplier will in fact supply in time and at appropriate quality, given the difficulties of coordination and implementation. A broad literature in organization theory has developed around the effort to map the various strategies of collaboration and control intended to improve the likelihood that the different components of the production process will deliver what they are supposed to: from early efforts at vertical integration, to relational contracting, pragmatic collaboration, or Toyota's fabled flexible specialization. The presence of a formalized enforceable contract, for outputs in which the supplier can claim and transfer a property right, may change the probability of the desired outcome, but not the fact that in entering its own contract with its clients, the company is making a prediction about the required availability of necessary inputs in time. When the company turns instead to the cloud of social production for its inputs, it is making a similar prediction. And, as with more engaged forms of relational contracting, pragmatic collaborations, or other models of iterated relations with coproducers, the company may engage with the social process in order to improve the probability that the required inputs will in fact be produced in time. In the case of companies like IBM or Red Hat, this means, at least partly, paying employees to participate in the open source development projects. But managing this relationship is tricky. The firms must do so without seeking to, or even seeming to seek to, take over the project; for to take over the project in order to steer it more "predictably" toward the firm's needs is to kill the goose that lays the golden

eggs. For IBM and more recently Nokia, supporting the social processes on which they rely has also meant contributing hundreds of patents to the Free Software Foundation, or openly licensing them to the software development community, so as to extend the protective umbrella created by these patents against suits by competitors. As the companies that adopt this strategic reorientation become more integrated into the peer-production process itself, the boundary of the firm becomes more porous. Participation in the discussions and governance of open source development projects creates new ambiguity as to where, in relation to what is "inside" and "outside" of the firm boundary, the social process is. In some cases, a firm may begin to provide utilities or platforms for the users whose outputs it then uses in its own products. The Open Source Development Group (OSDG), for example, provides platforms for Slashdot and SourceForge. In these cases, the notion that there are discrete "suppliers" and "consumers," and that each of these is clearly demarcated from the other and outside of the set of stable relations that form the inside of the firm becomes somewhat attenuated.

As firms have begun to experience these newly ambiguous relationships with individuals and social groups, they have come to wrestle with questions of leadership and coexistence. Businesses like IBM, or eBay, which uses peer production as a critical component of its business ecology - the peer reviewed system of creating trustworthiness, without which person-to-person transactions among individual strangers at a distance would be impossible - have to structure their relationship to the peer-production processes that they co-exist with in a helpful and non-threatening way. Sometimes, as we saw in the case of IBM's contributions to the social process, this may mean support without attempting to assume "leadership" of the project. Sometimes, as when peer production is integrated more directly into what is otherwise a commercially created and owned platform - as in the case of eBay - the relationship is more like that of a peer-production leader than of a commercial actor. Here, the critical and difficult point for business managers to accept is that bringing the peer-production community into the newly semi-porous boundary of the firm - taking those who used to be customers and turning them into participants in a process of coproduction - changes the relationship of the firm's managers and its users. Linden Labs, which runs Second Life, learned this in the context of the tax revolt described in chapter 3. Users cannot be ordered around like employees. Nor can they be simply advertised - to and manipulated, or even passively surveyed, like customers. To do that would be to lose the creative and generative social character that makes integration of peer production into a commercial business model so valuable for those businesses that adopt it. Instead, managers must be able to identify patterns that emerge in the community and inspire trust that they are correctly judging the patterns that are valuable from the perspective of the users, not only the enterprise, so that the users in fact coalesce around and extend these patterns.

The other quite basic change wrought by the emergence of social production, from the perspective of businesses, is a change in taste. Active users require and value new and different things than passive consumers did. The industrial information economy specialized in producing finished goods, like movies or music, to be consumed passively, and well-behaved appliances, like televisions, whose use was fully specified at the factory door. The emerging businesses of the networked information economy are focusing on serving the demand of active users for platforms and tools that are much more loosely designed, late-binding - that is, optimized only at the moment of use and not in advance - variable in their uses, and oriented toward providing users with new, flexible platforms for relationships. Personal computers, camera phones, audio and video editing software, and similar utilities are examples of tools whose value increases for users as they are enabled to explore new ways to be creative and productively engaged with others. In the network, we are beginning to see business models emerge to allow people to come together, like MeetUp, and to share annotations of Web pages they read, like del.icio.us, or photographs they took, like Flickr.

Services like Blogger and Technorati similarly provide platforms for the new social and cultural practices of personal journals, or the new modes of expression described in chapters 7 and 8.

The overarching point is that social production is reshaping the market conditions under which businesses operate. To some of the incumbents of the industrial information economy, the pressure from social production is experienced as pure threat. It is the clash between these incumbents and the new practices that was most widely reported in the media in the first five years of the twenty-first century, and that has driven much of policy making, legislation, and litigation in this area. But the much more fundamental effect on the business environment is that social production is changing the relationship of firms to individuals outside of them, and through this changing the strategies that firms internally are exploring. It is creating new sources of inputs, and new tastes and opportunities for outputs. Consumers are changing into users - more active and productive than the consumers of the industrial information economy. The change is reshaping the relationships necessary for business success, requiring closer integration of users into the process of production, both in inputs and outputs. It requires different leadership talents and foci. By the time of this writing, in 2005, these new opportunities and adaptations have begun to be seized upon as strategic advantages by some of the most successful companies working around the Internet and information technology, and increasingly now around information and cultural production more generally. Eric von Hippel's work has shown how the model of user innovation has been integrated into the business model of innovative firms even in sectors far removed from either the network or from information production - like designing kite-surfing equipment or mountain bikes. As businesses begin to do this, the platforms and tools for collaboration improve, the opportunities and salience of social production increases, and the political economy begins to shift. And as these firms and social processes coevolve, the dynamic accommodation they are developing provides us with an image of what the future stable interface between market-based businesses and the newly salient social production is likely to look like.

Notes

1. Richard M. Titmuss, *The Gift Relationship: From Human Blood to Social Policy* (New York: Vintage Books, 1971), 94.

2. Kenneth J. Arrow, "Gifts and Exchanges," *Philosophy & Public Affairs* 1 (1972): 343.

3. Bruno S. Frey, *Not Just for Money: An Economic Theory of Personal Motivation* (Brookfield, VT: Edward Elgar, 1997); Bruno S. Frey, *Inspiring Economics: Human Motivation in Political Economy* (Northampton, MA: Edward Elgar, 2001), 52-72. An excellent survey of this literature is Bruno S. Frey and Reto Jegen, "Motivation Crowding Theory," *Journal of Economic Surveys* 15, no. 5 (2001): 589. For a crystallization of the underlying psychological theory, see Edward L. Deci and Richard M. Ryan, *Intrinsic Motivation and Self-Determination in Human Behavior* (New York: Plenum, 1985).

4. Roland Bénabou and Jean Tirole, "Self-Confidence and Social Interactions" (working paper no. 7585, National Bureau of Economic Research, Cambridge, MA, March 2000).

5. Truman F. Bewley, "A Depressed Labor Market as Explained by Participants," *American Economic Review (Papers and Proceedings)* 85 (1995): 250, provides survey data about managers' beliefs about the effects of incentive contracts; Margit Osterloh and Bruno S. Frey, "Motivation, Knowledge Transfer, and Organizational Form," *Organization Science* 11 (2000): 538, provides evidence that employees with tacit knowledge communicate it to coworkers more efficiently without extrinsic

motivations, with the appropriate social motivations, than when money is offered for "teaching" their knowledge; Bruno S. Frey and Felix Oberholzer-Gee, "The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out," *American Economic Review* 87 (1997): 746; and Howard Kunreuther and Douslar Easterling, "Are Risk-Benefit Tradeoffs Possible in Siting Hazardous Facilities?" *American Economic Review (Papers and Proceedings)* 80 (1990): 252-286, describe empirical studies where communities became less willing to accept undesirable public facilities (Not in My Back Yard or NIMBY) when offered compensation, relative to when the arguments made were policy based on the common weal; Uri Gneezy and Aldo Rustichini, "A Fine Is a Price," *Journal of Legal Studies* 29 (2000): 1, found that introducing a fine for tardy pickup of kindergarten kids increased, rather than decreased, the tardiness of parents, and once the sense of social obligation was lost to the sense that it was "merely" a transaction, the parents continued to be late at pickup, even after the fine was removed.

6. James S. Coleman, "Social Capital in the Creation of Human Capital," *American Journal of Sociology* 94, supplement (1988): S95, S108. For important early contributions to this literature, see Mark Granovetter, "The Strength of Weak Ties," *American Journal of Sociology* 78 (1973): 1360; Mark Granovetter, *Getting a Job: A Study of Contacts and Careers* (Cambridge, MA: Harvard University Press, 1974); Yoram Ben-Porath, "The F-Connection: Families, Friends and Firms and the Organization of Exchange," *Population and Development Review* 6 (1980): 1.

7. Nan Lin, *Social Capital: A Theory of Social Structure and Action* (New York: Cambridge University Press, 2001), 150-151.

8. Steve Weber, *The Success of Open Source* (Cambridge, MA: Harvard University Press, 2004).

9. Maurice Godelier, *The Enigma of the Gift*, trans. Nora Scott (Chicago: University of Chicago Press, 1999), 5.

10. Godelier, *The Enigma*, 106.

11. In the legal literature, Robert Ellickson, *Order Without Law: How Neighbors Settle Disputes* (Cambridge, MA: Harvard University Press, 1991), is the locus classicus for showing how social norms can substitute for law. For a bibliography of the social norms literature outside of law, see Richard H. McAdams, "The Origin, Development, and Regulation of Norms," *Michigan Law Review* 96 (1997): 338n1, 339n2. Early contributions were: Edna Ullman-Margalit, *The Emergence of Norms* (Oxford: Clarendon Press, 1977); James Coleman, "Norms as Social Capital," in *Economic Imperialism: The Economic Approach Applied Outside the Field of Economics*, ed. Peter Bernholz and Gerard Radnitsky (New York: Paragon House Publishers, 1987), 133-155; Sally E. Merry, "Rethinking Gossip and Scandal," in *Toward a Theory of Social Control, Fundamentals*, ed. Donald Black (New York: Academic Press, 1984).

12. On policing, see Robert C. Ellickson, "Controlling Chronic Misconduct in City Spaces: Of Panhandlers, Skid Rows, and Public-Space Zoning," *Yale Law Journal* 105 (1996): 1165, 1194-1202; and Dan M. Kahan, "Between Economics and Sociology: The New Path of Deterrence," *Michigan Law Review* 95 (1997): 2477.

13. An early and broad claim in the name of commons in resources for communication and transportation, as well as human community building-like roads, canals, or social-gathering places-is Carol Rose, "The Comedy of the Commons: Custom, Commerce, and Inherently Public Property," *University Chicago Law Review* 53 (1986): 711. Condensing around the work of Elinor Ostrom, a more

narrowly defined literature developed over the course of the 1990s: Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (New York: Cambridge University Press, 1990). Another seminal study was James M. Acheson, *The Lobster Gangs of Maine* (New Hampshire: University Press of New England, 1988). A brief intellectual history of the study of common resource pools and common property regimes can be found in Charlotte Hess and Elinor Ostrom, "Ideas, Artifacts, Facilities, and Content: Information as a Common-Pool Resource," *Law & Contemporary Problems* 66 (2003): 111.