School of Economics and Finance

Occasional Papers

2005 Giblin Lecture

The Information Revolution and the Moral Economy of Innovation

Professor John Quiggin

Australian Research Council Federation Fellow, School of Economics
and School of Political Science and International Studies
University of Queensland

——— Occasional Paper No 8 ———

2005

School of Economics and Finance
University of Tasmania, Private Bag 85
Hobart Tasmania 7001

ISSN 1443-8593
ISBN 1 86295 300 7

School of Economics and Finance

Occasional Papers

2005 Giblin Lecture

The Information Revolution and the Moral Economy of Innovation

Professor John Quiggin[1]

Australian Research Council Federation Fellow, School of Economics and School of Political Science and International Studies

University of Queensland[2]

——— Occasional Paper No 8 ———

2005

School of Economics and Finance
University of Tasmania, Private Bag 85
Hobart Tasmania 7001

ISSN 1443-8593
ISBN 1 86295 300 7

---

[1] I thank Nancy Wallace for helpful comments and criticism. This research was supported by an Australian Research Council Federation Fellowship.

[2] EMAIL j.quiggin@uq.edu.au, PHONE + 61 7 3346 9646, FAX +61 7 3365 7299, http://www.uq.edu.au/economics/johnquiggin

Abstract

In the economy of the 21st century, economic and technical innovation is increasingly based on developments associated with the Internet and the World Wide Web. Unlike 20th century innovation, the most important developments in the Internet have been driven, not by research funded by governments and developed by corporations but by the collaborative interactions of Internet users. In most cases, innovation on the Internet has not been motivated by economic concerns or the prospect of profit. This raises the possibility of a world in which the leading sectors of the economy are driven by social interactions of various kinds, rather than by profit-oriented investment.

The Information Revolution and the Moral Economy of Innovation

Introduction

The observation that we are living through an 'information revolution' has become a cliché. Yet most discussion of this revolution takes it for granted that economic activity will go on much the same as before, only at a more frenetic pace. The archetypal metaphor is that of 'turbocapitalism' (Luttwak 1999).

However, the fundamental characteristic of information is that it is an almost perfect example of a pure public good. This point is captured in the slogan 'Information wants to be free'. Once produced, information can be used by as many people as can gain access to it, without reducing its availability to others. The crucial development of the Information Revolution has been to reduce the cost of distributing information, in many cases effectively to zero.

On the face of things, it seems implausible that an economic system based on a superbly efficient mechanism for the production and allocation of scarce resources, the price system, will continue as before in a situation where innovation is dominated by a free public good. On the contrary it is arguable that the most dynamic sectors of social activity are coming to be dominated by concerns other than resource allocation.

In this paper, the implications of the information revolution for the moral economy of innovation are examined.

The Internet and the Information Revolution

The expansion of the Internet has been one of the most spectacular technological and cultural developments humanity has ever seen. Twenty years ago, the Internet was an obscure network connecting a few thousand academics and students. Today it is a major feature of the lives of around a billion people, including the majority of the population in most developed countries.

Much discussion of the Internet is still shaped by the 'dotcom' boom of the late 1990s, in which as much as a trillion dollars was spent in the pursuit of largely spurious projects for commercial exploitation of the Internet. But

e-commerce and the dotcom boom were a diversion from the central path of Internet development, in which the crucial feature is the free sharing of information.

It was the free exchange of information that made the Internet so attractive and led to its success in competition both with proprietary online services such as Delphi and AOL and with more tightly-controlled government initiatives such as the X.25 protocol.

The free sharing of information was crucial both the design and construction of the Internet and to the provision of useful content. The software underlying Internet has always been based primarily on voluntary effort, and this has formed the basis of the open-source software movement, represented most prominently by Linux.

Before the rise of the WorldWideWeb the most interesting Internet content was found in the UseNet systems of newsgroups, covering almost every imaginable topic from atheism to zoology. Vigorous attempts to privatise and commercialise content were made during the dotcom era, with only modest success. Newspapers and other mass media built up substantial websites, attracting much attention away from Usenet and personal websites. However, attempts to charge for access to such sites have had very little success, and even reliance on advertising has proved to be problematic.

Since the end of the dotcom boom, the development of the Internet has relied increasingly on collaborative efforts based on free sharing of information.

Open-source software

The rise of the Linux operating system has focused attention on the idea of open-source software. The central idea of open-source is that software is distributed free of charge. Users can modify it as they see fit, and can choose whether to make their modifications publicly available, but cannot charge for the use of software derived from an open source program.

Open-source and the Internet have something of a symbiotic relationship. On the one hand, most of the software fundamental to the operation of the Internet is developed using some variation of the open-source model. In addition

to Linux, important examples include the Apache mail server

On the other hand, the Internet as a communications system makes it feasible to co-ordinate the efforts of hundreds or even thousands of programmers scattered around the world, using source code repositories such as SourceForge and Savannah.

Linux now contains millions of lines of code requiring thousands of programmer-years of development. Estimates of the value of this input for different distributions of Linux are between $US 1 billion and $US 2 billion (Wikipedia).

Blogs

A blog (this unfortunate term arose as a contraction of 'weblog' and has resisted periodic attempts to find a more appealing alternative) is simply a personal webpage in a journal format, with software that automatically puts new entries ('posts') at the top of the page, and shifts old entries to archives after a specified time, or when the number of posts becomes too large for convenient scrolling.

Some other elements, while not universal, are regarded by many as essential aspects of blogging. The most important are facilities for readers to make comments on individual posts and for other bloggers to link to posts with criticism, praise or merely to point to an interesting article. Most bloggers also locate themselves within a larger community through the device of a 'blogroll', that is, a sidebar with a list of permanent links to other blogs likely to be of interest to readers.

Although blogs can be run either by individuals or groups, content is typically individual. That is, each post is written by a single author, as is each comment. There is of course, nothing to prevent joint composition of either posts or comments, but standard blog software does nothing to facilitate such collaboration, and in some ways actively discourages it. For example, the standard setup for a group blog does not allow members of the group to edit each others' posts.

Blogs serve many purposes, from online diaries to corporate public

relations. But the bloggers who have attracted the most public interest are those engaged in political and cultural debate. From tiny beginnings in the late 1990s, blogging has grown at an exponential rate. An indication is provided by the free service Technorati which monitors links between blogs (links, rather than monetary flows are the main currency of the blog world). In September 2003, Technorati claimed to watch nearly 4 million weblogs, with more than half a billion links. By October 2005, the number had risen to 20 million. Other estimates are as high as 50 million, but these are inflated by various online communities that automatically create blog-style diaries for new members and by the use of spurious blogs to manipulate the rankings of search engines like Google (a variant of 'spam').

Wikis

By contrast with blogs, and most other websites, wikis are designed to facilitate editing by as many people as possible. The most prominent single example of a wiki is Wikipedia http://en.wikipedia.org , an online encyclopedia, from which information for this section has been drawn.

A wiki enables documents to be written collectively (co-authoring) in a simple markup language using a web browser. A wiki is a collection of pages, which are usually highly interconnected via hyperlinks,—in effect, a very simple relational database. The name was based on the Hawaiian term wiki wiki, meaning "quick" or "informal."

Wikis generally practice the philosophy of making it easy to correct mistakes, rather than making it difficult to make them. A typical wiki allows editing by anyone visiting the relevant page, without any requirement for membership of a specified group, or even for registration.

While wikis are very open, they provide a means to verify the validity of recent additions to the body of pages. The most prominent, on almost every wiki, is the "Recent Changes" page—a specific list numbering recent edits, or a list of all the edits made within a given timeframe.

Most of the relevant features of wikis may be illustrated by Wikipedia, which is summarised by itself, as follows:

> Wikipedia is a Web-based, free-content encyclopedia written collaboratively by volunteers and sponsored by the non-profit Wikimedia Foundation. It has editions in roughly 200 different languages (about 100 of which are active) and contains entries both on traditional encyclopedic topics and on almanac, gazetteer, and current events topics. Its purpose is to create and distribute a free international encyclopedia in as many languages as possible. Wikipedia is one of the most popular reference sites on the internet,[1] (http://en.wikipedia.org/wiki/Wikipedia#endnote_PopularityRef) receiving around 60 million hits per day.
>
> Wikipedia contains approximately 1.6 million articles. More than 600,000 of these are in English, more than 240,000 in German, and more than 100,000 each in Japanese and French. It began as a complement to the expert-written Nupedia on January 15, 2001. Having steadily risen in popularity, [2] (http://en.wikipedia.org/wiki/Wikipedia#endnote_Popularity) it has spawned several sister projects, such as Wiktionary, Wikibooks, and Wikinews. It is edited by volunteers in wiki fashion, meaning articles are subject to change by nearly anyone. Wikipedia's volunteers enforce a policy of "neutral point of view" whereby views presented by notable persons or literature are summarized without an attempt to determine an objective truth. Because of its open nature, vandalism and inaccuracy are problems in Wikipedia.
>
> The status of Wikipedia as a reference work has been controversial, and it is both praised for its free distribution, free editing and wide range of topics and criticized for alleged systemic biases, deficiencies in some topics, and lack of accountability and authority when compared with traditional encyclopedias. Its articles have been cited by the mass media and academia and are available under the GNU Free Documentation License. Its German language edition has been distributed on compact discs, and many of its other editions are mirrored or have been forked by websites.

Web 2.0

Open-source software, blogs and wikis are just the most prominent of a string of web-based initiatives relying primarily on collaboration between users. These developments, taken as a group, have been referred to by some commentators as ‘Web 2.0’ (O’Reilly 2005)

Motives for contributing

There are all sorts of motives which might lead people to contribute to networked social capital, for example by participating in various aspects of blogging (make posts and comments, linking and blogrolling, improving software, various kinds of metablogging). Possible motives include altruism, self-expression, advocacy of particular political or social views, display of technical expertise, social interaction and so on.

In general, these motives are complementary or at least mutually consistent. However, motives like these do not co-exist well with a profit motive. Benkler notes the absence of monetary side payments in the case of car-pooling and this is typical of co-operative endeavours of various kinds. There is a strong desire to keep money out of this kind of activity, except where there is necessary interaction with the monetary economy (for example, car-poolers contributing to tolls, bloggers sharing the costs of hosting services and so on).

Why is this? At a superficial level, it's obvious that people act differently, and are expected to act differently, in the context of relationships mediated by money than in other contexts. Behavior that would be regarded favorably in a non-monetary context is regarded as foolish or even reprehensible in a monetary context.

One of the most important general differences relates to rationality and reciprocity. In a non-market context, careful calculation of costs and benefits and an insistence on exact reciprocity is generally deprecated. By contrast, in market contexts, the first rule is never to give more than you get. (As economists hasten to point out, the existence of gains from trade means that this does not imply conflict between the parties to a market transaction: both can gain more than they get. On the other hand, this doesn't mean there isn't conflict: both could do better at the expense of the other).

Why is it more important to observe this rule in market contexts. One reason is that markets create opportunities for systematic arbitrage that don't apply in other contexts. In an environment where trust is taken for granted, a trader who consistently gives slightly short weight can amass substantial profits. If trading partners assume honourable behavior, none will suffer enough to notice.

This is much more difficult to do in ordinary social contexts.

Similar points can be made about other motives. There are a whole range of sales tricks designed to exploit altruism, friendship, desire for self-expression and so on. Hence, to prosper in a market context, it is necessary to adopt a view that 'business is business', and to (consciously or otherwise) adopt a role as a participant in the market economy that is quite distinct from what might be conceived as one's 'real self'.

At this point in my argument, it seems clear that sociology has a lot to contribute, obvious instances including Weber, Goffman and Hirschman. Most importantly, there seems to be some potential for insight into the question of the circumstances in which personality is role-specific in some fundamental sense, rather than in the trivial sense in which a role is defined by the performance of specific functions.

Many of the same points apply in the context of bureaucratic rationality, though here it is the interests of the other party that are being protected. When dealing with a representative of a bureaucracy, we (mostly) expect consistent application of rules, rather than an adherence to standard social norms such as 'look after your family/mates before others".

The process of innovation

Households, markets and governments

In broad terms, the economic structure of the 20th century economy[3] may be described in terms of three sectors: household, market and government. In this model, economically relevant activity takes place primarily in the market and government sectors.

Market sector final output is sold to households, and constitutes around 80 per cent of final consumption. Households supply labour and other inputs to market sector production and receive flows of wages, rent and dividends in

3 This structure began to emerge in the second half of the 19th century and was fully developed with the rise of the welfare state after World War II. The modern economy displaced an agriculture-based economy, in which market exchange played a secondary role.

return.[4]

A crucial element of markets in the 20th century economy is reliance on the impersonal rationality of the price mechanism. Firms provide goods and services, not out of goodwill to consumers or a desire to display excellence (although such motives are frequently referred to in corporate advertising), but because they can provide the goods and services at a price acceptable to consumers and make a profit in the process.

Governments provide services that, for one reason or another, cannot be adequately delivered by the market, at least in the judgement of the society concerned. These typically include health, education and a range of 'public goods'.

While the category of 'public good' is rather vague in practice, economists have modelled public goods in terms of two crucial characteristics. The first is nonrivalry, meaning that use of a public good by one person does not diminish its availability for others. The second is nonexcludability, meaning that once the good is produced it is impossible to prevent people from using it, and therefore impossible to charge a market price for use.

The paradigm case of a pure public good is information. My possession of information does not stop you from using it. Moreover, once knowledge is public it is difficult impossible to prevent anyone from gaining access to it, though it may be possible to restrict or prohibit particular uses of information.

The household sector plays a fairly minor role in the modern economic model. Economists such as Becker (1976) have attempted to apply a market-based analysis to the use and distribution of resources within households., using such tools as shadow prices and marginal rates of substitution, but with little lasting success. Somewhat more progress has been made using models of bargaining, but the household remains, in large measure, outside the sphere of economic activity as we normally understand it.

4 In this context, the household sector is defined broadly to include everything done by individuals and families that is not mediated through markets or governments.

Innovation and productivity growth

The standard process of innovation in the 20th century model has three stages. Pure or fundamental research, producing new knowledge as a public good, is undertaken by the public sector, in universities and other research institutions. Applied research and development is undertaken in the market sector, resulting in the creation of new products. These products are supplied to the household sector where they are passively consumed.

A slightly more complex version of the model takes account of process innovations, which enhance productivity in the market sector, and of innovations leading to new capital goods, also used in market sector production. The resulting gains in productivity are referred to as 'embodied' when they involve new kinds of capital goods and 'disembodied' when they involve process improvements.

In principle, these process and capital innovations might apply to the government sector. However, because of the difficulty of measuring government-sector output in market terms, the standard practice of national accounts in the 20th century was to treat the value of government sector output as being equal to the value of inputs, implicitly defining the rate of productivity growth in the government sector to be equal to zero.

Improvements in productivity are, ultimately, the only source of sustainable economic growth. In the 20th century model, (measured) improvements in productivity growth took place only in the market sector. Thus, the market sector was central to economic growth and economic analysis.

The Internet model of productivity growth

The model of productivity growth associated with the Internet is radically different from the 20th century model described above. The reliance of Internet innovation on collaborative processes reverses the flow of conservation. Instead of being a passive recipient of innovations developed elsewhere, the household sector, broadly defined to include all the activities undertaken by individuals other than for pay or to produce items that can be sold or licensed, is now the focus of some of the most significant innovations that are taking place.

This is not a wholly new development. The idea that innovation is driven, to a large extent, by customers, has been explored by a number of analysts, notably von Hippel (1988). However, von Hippel focuses primarily on large industrial customers of companies supplying capital goods. The implications of innovation undertaken by thousands of individual users of software-based products are more far-reaching and profound.

Blogs provide one of the most important focal points for innovation in the Internet. A clearer example of technological innovation driven by blogs is that of the various versions of RSS/RDF. Cutting through the tangle of acronyms, RSS is an effective implementation of what used to be called 'push technology', The idea is that, instead of visiting the sites on your bookmark list in turn to see whether there is anything new, the sites themselves generate a short report noting recent posts and comments. A program such as NetNewsWire, called an aggregator, collects these reports for all the sites you wish to visit.

As with many innovations, there is a need to generate a virtuous circle. The success of the technology requires a workable underlying RSS/RDF specification, user-friendly aggregator software, and a large enough number of sites generating RSS feeds to sustain user interest. Blogs and their readers provided the third component of this cycle. The time-consuming nature of blogging meant that blog readers welcomed the RSS innovation, even with the rough edges that characterised its early stages. This in turn produced pressure on bloggers to implement RSS and therefore on developers of blog software to make this as easy as possible.

Once the groundwork had been laid in the blogosphere, RSS spread rapidly. Most major news sites now offer RSS feeds, and web browsers, such as Safari, are beginning to offer RSS as a routine part of their service.

Going beyond its direct importance, RSS/RDF is the most prominent single example of the benefits of the XML/XHTML approach over the simpler HTML which served the World Wide Web for its first decade. Progress in Web and Internet technology is driven by the most demanding users, and, increasingly, these are the members of creative communities such as the blogosphere.

Moore's Law

In assessing the significance of collaborative innovation, it is important to observe the whole process is underpinned by steady improvements in the speed of computation and communication, and these largely follow the traditional model, with larger and larger capital investments being required for each generation of computer chips. However, without the Internet, these improvements would have produced little more than faster word processors and spreadsheets.

There is no sign that the rate of progress in computer technology is slowing down noticeably. A doubling time of two years for chip speed, memory capacity and so on implies a thousand-fold increase over twenty years. There are two very different things this could mean. One is that computers in twenty years time will do mostly the same things as at present, but very fast and at almost zero cost. The other is that digital technologies will displace analog for a steadily growing proportion of productive activity, in both the economy and the household sector, as has already happened with communications, photography, music and so on. Once that transition is made these sectors share the rapid growth of the Internet sector.

In the first case, the contribution of computer technology to economic growth gradually declines to zero, as computing services become an effectively free good, and the rest of the economy continues as usual. Since productivity growth outside the sectors affected by computers has been slowing down for decades, the likely outcome is something close to a stationary equilibrium for the economy as a whole.

But in the second case, the rate of growth for a steadily expanding proportion of the economy accelerates to the pace dictated by Moore's Law. Again, communications provides an illustration after decades of steady productivity growth at 4 or 5 per cent a year, the rate of technical progress jumped to 70 per cent a year around 1990, at least for those types of communication that can be digitized (the move from 2400-baud modems to megabit broadband in the space of 15 years illustrates this).

One way of expressing this dichotomy is in terms of the aggregate elasticity of demand for computation. If it's greater than one, the share of computing in the economy, expressed in value terms, rises steadily as computing

gets cheaper. If it is less than one, the share falls. It is only if the elasticity is very close to one that we continue on the path of the last couple of decades, with continuing growth at a rate of around 3 per cent.

This kind of result, where only a single value of a key parameter is consistent with stable growth, is sometimes called a knife-edge. Reasoning like this can be tricky - maybe there are good reasons why the elasticity of demand for computation should be very close to one. One reason this might be so is if most problems eventually reach a point, similar to that of weather forecasting, where linear improvements in performance require exponential growth in computation.

If the solution to a problem involves components that are exponential (or worse) in complexity, initial progress may be rapid as linear or polynomial components of the problem are solved, but progress with the exponential component will at best be linear, even if the cost of computation is itself declining exponentially.

So far it seems as if the elasticity of demand for computation is a bit greater than one, but not a lot. The share of IT in total investment has risen significantly, but the share of the economy driven primarily by IT remains small. In addition, non-economic activity like blogging has expanded rapidly, but also remains small. The whole thing could easily bog down in an economy-wide version of 'Intel giveth and Microsoft taketh away'.

A post-economic society?

Discussion of the Internet is often Utopian. Indeed, the Internet is rather like Harry Potter's Mirror of Erised, which shows the viewer whatever they most want to see. Among the academics who built the Internet this was a co-operative world in which sharing based on mutual esteem would displace the profit motive and render large corporations obsolete. When the emergence of the World Wide Web coincided with the 1990s stock market boom, the mirror showed a route to instant riches, and produced the dotcom mania.

Now something more akin to the older vision of collaborative creation is re-emerging. Ideas of this kind range from the relatively sober (Lessig 2001) to the speculative (O'Reilly 2005) to the positively visionary (Kevin Kelly 2005). At the

extremes, this kind of utopianism merges into speculation about a 'Singularity' that will fundamentally transform human existence (Kurzweil 2005).

Focusing on the realistic end of the spectrum, the steady growth in the volume and sophistication of highly hypertextual components of the Web, such as blogs and wikis is giving rise to a steady flow of innovations that are fundamentally transforming the way in which people interact with information. As yet, the impact of these developments is confined primarily to early adopters in the household sector, though these already number in the millions. However, it is already clear that where the early adopters have led, the mass of household users will eventually follow.

Given this broad base, it is safe to predict that the hypertextual Web will drive productivity innovations in the business and government sector for some years to come. The crucial question is whether this kind of model can be extended to broader ranges of social and economic activity, particularly those involving interaction with physical goods. The ways in which this might happen remain unclear. One possibility is through the extension of open and programmable intelligence to household goods and perhaps to the way in which houses themselves work. An early example is the TiVo programmable service, the capacities of which have been extended by 'hackers' apparently with the tacit compliance of the supply company.

Going beyond this, it is possible to foresee a revival in the long-nascent field of home automation, driven by collaborative software development. Although protocols such as X.10 have existed for many years, they have not been successful as a basis for commercial products. Perhaps collaborative Internet development may be able to succeed in this area.

Economic and policy implications

First, to the extent that innovation and productive growth arise from activities that are pursued primarily on the basis non-economic motives, the link between incentives and outcomes is weakened. This in turn undermines the reationale for policies aimed at sharpening incentives and ensuring that everyone engaged in the production of goods and services is exposed to the incentives

generated by a competitive market. Such policies represent the core program of 'economic rationalism', the set of ideas that dominated Australian public policy in the 1980s and 1990s.

The central focus of economic rationalism was 'getting prices right' by removing tariffs, subsidies and other distortions and by exposing public sector activities to the market competition. These policies yielded important benefits. Looking at the rise of the Internet in Australia, however, it is easy to see that economic rationalism had nothing to do with it. The Australian component of the Internet (like its overseas counterparts) was developed within the university sector with a combination of small amounts of funding for bodies like AARRNET and large inputs of unpaid labour (either academics' spare time, or a diversion of effort from their official tasks, depending on how you want to look at it). When universities finally got into commercialisation of the Internet, as with the float of Melbourne IT, the result was an unedifying boondoggle, with no net contribution to the public good.

Similar points may be made in relation to another core belief underlying the the reforms of the 1980s and 1990s, the idea that the 'cold shower' of competition drives productive innovation. Again, there is nothing in the rise of the Internet to support this idea. Indeed, it is hard to imagine that, in the post-reform Australian university sector of the late 1990s, anyone would have had time for something as apparently frivolous as an all-purpose electronic communications network — the struggle for survival focused attention on the short term.

The point may be made more broadly in terms of a distinction between rationality and creativity. The 20th century model of innovation is one dominated by rationality. Creativity is needed for the initial generation of ideas for pure research, but plays little role thereafter. In the 21st century model, the roles are, if anything, reversed. Rational design is still fundamental in keeping the whole system of the Internet going, but the day-to-day process of innovation is driven by creative discovery and exploitation of the possibilities of the new medium.

This distinction points up a potential role for public policy. One of the difficulties in a system of innovation based on spontaneous contributions is that the allocation of effort to tasks will mainly reflect characteristics other than the

benefits flowing from its completion. In open-source software for example, programmers may be keen to work on tasks that are technically challenging, or will bring them fame or will solve a significant problem for some particular group of people they want to help. They will be less likely to work on boring tasks that improve the functioning of the system but don't have a big impact on anyone in particular.

Similarly, in examining a collectively-produced product like Wikipedia it is easy to find passages and whole pages that are poorly-written and in need of editing. But contributors mostly find this an unrewarding task compared to writing new pages, adding new facts or correcting factual errors. Quite possibly the only person who will actually notice such a contribution is the original author and they may not be grateful. The benefit to readers from a better article may not provide sufficient motivation. Public policy could assist in filling this gap, either through direct provision of undersupplied Internet services or through the provision of funding to groups undertaking such provision, such as the Wikipedia foundation.

On the other hand, the analysis presented here casts doubt on the currently popular policy of encouraging and subsidising venture capital enterprises. The investments encouraged by venture capitalists in the dotcom era may have rewarded their promoters, but they produced little of lasting social value, at least by comparison with the vast sums that were invested.

In summary, rather than seeking to drive people harder in the search for increased productivity, government policy should be oriented towards making room for creativity and facilitating its expression. Similarly, while competition has its place, public policy should be at least as much concerned with promoting co-operation.

References

Becker, G. (1976) *The Economic Approach to Human Behaviour,* University of Chicago Press, Chicago.

Kelly, K. (2005), 'We Are the Web', Wired, http://www.wired.com/wired/archive/13.08/tech.html?pg=2&topic=tech&topic_set=,

Kurzweil, R. (2005) *The Singularity is Near,* Viking, New York.

Lessig, L. (1999) *Code and Other Laws of Cyberspace,* Basic Books, New York.

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

Luttwak, E. (1999) *Turbo Capitalism: Winners and Losers in the Global Economy,* Harper Collins, New York.

O'Reilly, T. (2005), 'What Is Web 2.0?: Design Patterns and Business Models for the Next Generation of Software', http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html,

von Hippel, E. (1988) *The Sources of Innovation,* Oxford University Press, Oxford.