Prometheus, Vol. 27, No. 4, December 2009

Routledge
Taylor & Francis Group

Information Technology and the Economics of Storing, Spreading and Generating Knowledge

PETER E. EARL

Abstract *This paper examines the future for the growth and transmission of knowledge with a particular focus on the collision between the zero-marginal-cost economics of online publishing and the finite attention spans of those who download content. The paper tries to show that while it is likely that the information revolution will produce much greater equality of access to knowledge, its impacts on the efficiency with which people pick up capabilities or generate new knowledge remain debatable.*

Keywords: open access; zero-marginal-cost; online publishing; cognitive dissonance; price discrimination; attention

Introduction

This paper uses inputs from information economics and behavioural economics to examine how information technology changes the processes by which knowledge is created and transmitted, with a focus on changes in both cost and quality. It is divided up in a sequence that begins with access to knowledge resources and ends with the durability of electronic knowledge resources after considering how information technology affects the acquisition and creation of knowledge. Firstly, we examine the economics of electronic access to knowledge resources and the long-term implications of information technology for the infrastructure of universities and the strategies of academic publishers. Secondly, we use ideas from both information economics and behavioural economics to consider how easy electronic access to knowledge has different impacts on the development of know-how from those that it has on improving knowledge of facts. Thirdly, we challenge the view that information technology necessarily helps researchers to generate creative outputs at lower cost than in past eras where word processing technologies were not in use. Fourthly, prior to some concluding comments, we consider how perishable knowledge is when it is stored electronically.

Prometheus ISSN 0810-9028 print/ISSN 1470-1030 online
http://www.informaworld.com/journals
DOI: 10.1080/08109020903414275

Openness of Access

The switch to online journals and books removes the need for academic staff and students to purchase personal copies as a means of ensuring access. The need to possess a hard copy is also fading fast. To be sure, reading on-screen is not the same as reading from a hard copy but, as the success of Amazon.com's Kindle pdf-players indicates, things are improving rapidly. Future generations of academics and students will store their personal collections of books and articles in portable devices whose screens suffer no significant disadvantages (except perhaps tactile ones) compared with a traditional printed volume. Just as has happened with music via the rise of MP3 music files and iPOD devices, so we shall be able to have all of our knowledge resources at hand at any time, with no access problems associated with limited library opening hours or someone else having borrowed the work we wish to use. Already some universities are offering lecture presentations for iPOD use via iTunes U and with the collapse in price of portable 1TB hard drives and spread of WiFi broadband access the time is ripe for the convergence of these technologies in devices no bigger than a typical novel or monograph but which enable their owners to access, store and play back huge amounts of music, video and document data. These devices will doubtless offer far superior reading functionality than the traditional book: for example, our pdf players will be able to display text at a size that suits our eyes and will be able to present text in digital audio form for us to listen to whilst on the move. They will also offer far better search capabilities than traditional indexes and the ability to download items from reference lists.

Such products will have significant implications for university managers. In the past, academics who took their jobs seriously spent large sums on their personal libraries and had to incur the time costs of ordering books and journals and justifying their claims for them when filling out their income tax returns. A free personal supply of a much-used journal was an inducement to agree to serve as an editorial board member, in addition to the status such membership might bring. In future, academic salaries will no longer need to compensate for this where academics work at institutions whose libraries are so well resourced that they have no need to buy books and journals for their personal libraries. (Amazon.com begs to differ on this scenario as its model involves downloading books from its store to Kindle devices, rather than downloading library e-books.) Bigger savings will come in terms of physical infrastructure: when access to library holdings is non-rival and downloaded items are stored in readers' electronic devices, the academic offices to which we have become accustomed will no longer be necessary. As buildings get refitted to accommodate more staff, or new university buildings are erected, the work environment in universities will come increasingly to resemble that of the corporate world, with much smaller private offices that offer peaceful working conditions, and shared rooms to be used for meetings. There may also be some implications for staff mobility: no longer will the dread of packing and unpacking one's office loom large if one is presented with a chance to move on.

Another likely change is in relative access to published work among institutions of differing wealth and age. The marginal costs of providing access to e-books and journals is close to zero, both for new works and for past hard-copy publications that have been scanned into electronic form. Similarly, the receiving institution will not be constrained by shortages of physical storage space and the marginal cost of server capacity is trivial. The only constraint a small, new university has as a

purchaser of e-books and journals is its budget. Hence we should expect the publishers to develop very sophisticated price discrimination strategies that offer download access at prices that differ considerably between institutions, based on ability and willingness to pay. The cost of a bundle of journals would thus depend on negotiations that focused on the numbers of students and researchers, the mix of teaching and research across discipline areas and specialization within them, and income per student. As a very simple example, if universities A and B enjoy equal funding per head and a similar range of discipline areas but A has twice as many students as B, then we might expect A's subscription to a particular set of journals to cost twice that of B, resulting in identical costs per student, whereas in the past both institution might have had to pay the same price for each annual subscription with the result that A's library held twice the range of titles. It is possible that such a vision has been a major driver of the mergers and acquisitions that have taken place in academic publishing in the last couple of decades. Command over a great portfolio of journals offers bundling economies in terms of the potential for transaction cost savings associated with negotiating the one big deal with each library but, where multiple deals are done, it also permits block-booking economies akin to those enjoyed by diversified media firms when they sell advertising space.[1]

Publishers need to be able to prevent leakage of their products between institutions that are being charged different prices if such price discrimination strategies are to result in a great levelling of access to knowledge, both within and between nations. A small university in, say, Africa must not be able to download an entire collection and resell it to a large university in a rich country, and neither should it be possible for an individual to download a collection and on-sell it at a price that undercuts the price the publishers are demanding from other institutions. The means for doing this are already in place in the form of software to detect wholesale downloading of a publisher's products via automated web-crawlers and to prevent complete volumes of journals from being downloaded in a single session. This technology enables publishers to impose substantial costs on university libraries who mistakenly grant access rights to devious individuals, such as sabbatical visitors from much poorer institutions: if the publishers detect abnormal downloading patterns they can shut off all access via that library's portal until the library finds and excludes the culprit (after requiring deletion of the article files from the mass download).

The zero-marginal-cost economics of online publishing also imply an alternative scenario by which this levelling of access comes about, but one which has rather different consequences for those who produce the knowledge. Here, academic publishers make their products freely downloadable with their fixed costs and profits being covered by charges levied on the authors of the research, or by much more substantial submission fees than some journals currently charge. If these fees were uniform for a particular journal the benefits in terms of universal access to the eventual publications in the journal would be somewhat tempered by the greatly increased inequity in access to being published. However, price discrimination could be employed here, too, with different charges based on the academic's geographical zone and status.

The second scenario is perhaps the more probable one since, by equating the access cost to the zero marginal cost of production and delivery, it removes any incentive for users to devise means of getting round access problems caused by price discrimination being less than perfect, particularly en route to the end state

envisaged in the first scenario. So long as access to published works is fee-based and differs between institutions, a world of knowledge increasingly dominated by pdf files rather than hard copies becomes increasingly ripe for the emergence of web-based file sharing systems that do for pdf files what Napster did for MP3 audio files. While we might imagine that, say, Science Direct would be likely to get as aggressive with participants in such a system as Metallica got with Napster, we might also note that the ensuing bankruptcy of Napster has not eliminated peer-to-peer file exchange systems for recorded music; indeed, aided by BitTorrent technology, such sharing now takes place on a far greater scale than it did via Napster. Even without having to contend with such a technology, commercial publishers that make contributions to knowledge available electronically are constrained in their ability to enforce copyright by a crucial difference between publishing in the academic world and in the audiovisual entertainment industry. In the latter, the artists who produce the creative inputs normally share with the publisher an interest in ensuring that those who obtain access do so by paying for it. A performer or writer whose income is derived from sales of their creative work will not respond to requests that they provide free copies unless this will eventually lead to higher revenues from their charged-for products. By contrast, if an academic receives an e-mail asking for a pdf of work that the publisher prevents them from making downloadable from their home page, they are likely to respond positively. They do not reduce their income by giving away access to their work and may even ultimately benefit to the extent that doing so results in higher rates of citation. Replying to an e-mail that requests a pdf of an article takes seconds, unlike making a physical copy and putting it in the traditional mail.

So far, however, we are a considerable way from either equal access scenario. Some academic publishers, such as Springer, are hedging their bets and trying a hybrid system in which they both charge for electronic access to their journals and give authors of articles the chance to buy universal reader access to their papers for a non-trivial fee (roughly the cost of an international conference trip, in terms of opportunity cost of the academic's research funds). At present, rational authors probably opt for the former strategy and hope to receive e-mail requests from those who discover their papers but whose libraries do not subscribe to the journals in which they are published. However, to the extent that this minor hurdle for would-be readers is a hurdle nonetheless, those academics who can afford to pay the universal access fee will tend to garner more citations (as well as enjoying the status benefits of being able to demonstrate publicly that they or their institutions can afford to pay), so in the long run it will be advantageous to follow suit. As more and more academics opt to pay to permit universal access the publisher will run into a complex trade-off. The more articles that appear with universal access, the more revenue the publisher earns from authors but the less attractive the journal is becoming to libraries. Hence subscription revenues will fall unless a consequent rise in the journal's impact factor due to its wider readership enables it to attract an offsetting increase in the quality of papers that are received from authors who are not prepared to pay for universal access.

In the current transitional state, some kinds of users of published knowledge will be experiencing diminished or more expensive access than they would have enjoyed in the world of hard copy publishing. The demise of hard-copy journals and books poses a potentially serious difficulty for retired academics who wish to continue doing research but whose former employers do not offer them some kind

of honorary position that provides them with ongoing access to online material in their libraries. Though the marginal cost to the former employer of allowing such access may be close to zero, university libraries may be prevented by the publishers from allowing such access to anyone other than current staff and students. Electronically stored knowledge is open to such restrictions in a way that hard-copy knowledge it not: publishers cannot readily monitor who walks into a library, whereas they can readily restrict downloads of their products via a library's Internet portal.

This problem may blight the retirement dreams of academics who had hoped that the information technology revolution might make it possible for them to enjoy several decades of research after retiring to locations remote from their former places of work, something that would have seemed much harder if they had to rely on hard copy library materials. If we cannot rely upon universities subverting the greed of publishers by freely awarding 'honorary research associate' kinds of roles and keeping their retired staff listed as if they are still employees, we might also be concerned about the likely unwillingness of publishers to extend their price discrimination strategies to include concessionary rates for retired academics to gain electronic access to their databases: someone applying via a non-university e-mail account could actually be on a mission to profit by downloading their material and on-selling it to others. Whereas a hard-copy journal that has been obtained on a concessionary subscription can only be on-sold once, electronic copies can be repeatedly on-sold and hence it is far more worthwhile for would-be pirates to incur the transaction costs of making this happen. However, even if retired academics find themselves facing the inconvenience and cost of purchasing individual journal article downloads, in contrast to the ability of current university staff to browse freely across a great range of electronic resources, we should not exaggerate their plight: they will normally still be able to view abstracts freely and find author e-mail addresses as a means to a good chance of getting electronic access rapidly and with no financial outlay.

The Facilitation of Improved 'Know-How' Rather than Merely 'Know-That'

Information obtained online, like that retrieved from traditional systems, may be inaccurate or at odds with other information. Although access to greater volumes of information may thus require greater use of capabilities for separating knowledge from quackery, the information technology revolution seems unambiguously to have reduced the costs of obtaining knowledge of the 'know-that' kind.[2] For example, now, if we watch a movie and end up arguing about which other movies one of the actors appeared in, or whether so-and-so is still alive and active, we can resolve the issue in a few moments via Wikipedia, where facts are updated with remarkable rapidity. Such pieces of 'trivia' would have been very costly to acquire or verify without this technology and the ability 'to Google': a trip to one's local library in the past would, at a much higher cost, have provided less of a prospect of an accurate result. However, for knowledge of the 'know-how' kind—i.e. knowledge that is necessary to achieve a capability to perform a particular activity—the information revolution earns a rather more patchy scorecard: it has enormous potential to spread know-how widely and at a fraction of the cost of traditional channels but there is also the risk that in some contexts it will work in a counterproductive manner, with growth in what is merely 'know-that' giving an illusion of growth in 'know-how'.

For example, on the positive side, let us consider the impact of the information revolution on processes by which the know-how required to play musical instruments is acquired. Here, its most revolutionary consequences lie with popular musical expertise rather than in the further democratization of the ability to learn how to play the classical repertoire on classical instruments. Prior to the advent of sound and video recordings, musicians such as Paganini, who devised new techniques and composed work that pushed the boundaries of what was technically possible, could delay the spread of their know-how and ensure high demand for tickets to their performances by refusing to provide master-classes and not publishing their compositions in printed form. Now, however, given the ready access to classical music in recorded and sheet form (with fingerings indicated) and the established institutions for passing on classical techniques, it is doubtful that the information revolution is likely to affect the levels of attainment of those who are willing to make serious attempts to become competent classical musicians. The knowledge necessary to achieve competence as a classical musician has long been accessed by a combination of oral instruction and printed sheet music. A graded system of examinations provides a source of motivation to practise and to try to absorb advice from one's teacher as well as a means of certifying what level of competence has been achieve.

Matters are very different in the case of popular music, where tuition is less readily available in schools and where many musicians tend to be self-taught. For budding guitarists seeking to learn how to play particular pieces of recorded popular music the typical problem used to be the absence of comprehensive and accurate sheet music transcriptions, or of *any* sheet music at all, even for many recordings that had enjoyed huge commercial success. It was thus necessary to try to pick up what was happening directly from the recording, by repeatedly playing sections of a vinyl disc or tape cassette. Slowing down a recording to work out what was going on was very difficult. Matters improved during the 1980s and 1990s with the advent of video recordings and spread of home VCRs that enabled televised performances to be captured and studied. Firms such as Hal Leonard Music began to publish a limited range of note-for-note transcription albums and sometimes complete scores of particular albums, though these remained far more expensive than classical sheet music. A–B repeat facilities on CD players made it much easier to study small portions of recorded music. However, with the rise of the Internet, there has been a dramatic increase in the ability to find out how virtuoso performers of popular music do what they do, and what notes they are playing, combined with a collapse in the cost of obtaining this knowledge. This is largely a result of musicians sharing the investments they have made in acquiring different pieces of musical knowledge, such as YouTube video clips of established performers including many that are impossible to obtain commercially, or clips of themselves demonstrating particular techniques, or transcriptions they have made of their favourite recordings. Music tablature websites such as Protabs or Guitar-Pro charge modest sums for software that enables free downloading of, in many cases, complete scores of transcriptions of popular music recordings, along with suggested fingerings. Moreover, these include realistic digital audio versions of the instruments that can be remixed to enable individual instruments to be focused upon or backing tracks to be created, along with tempo flexibility and A–B looping to permit short passages to be studied carefully. Although the kinds of comments posted on YouTube music video clip pages are often far from civil, some do include

remarks that amount to instructions aimed at helping viewers unpack the tacit knowledge that is being employed by the performer in the clips.

The Internet has thus provided a low-cost means of accessing the know-how required by those who wish to play non-classical musical instruments to a very high standard. A convergence is in process between what can be done by these kinds of musicians and by those who have been formally trained by traditional methods. There has been an explosion in the number of gifted amateurs, whose achievements can inspire and be copied by others. For example, on YouTube one can see Paganini's Moto Perpetuo and 5th Caprice being performed not only by some of the greatest violin players but also by electric guitarists both known and unknown, or implausible achievements such as Rimsky-Korsakov's 'Flight of the Bumblebee' being played on electric bass. With exemplars such as these available, along with tablature sites, both aspirations and attainments can increase.

In principle, we might hope that the information revolution would produce rather similar kinds of effects in other areas of knowledge acquisition. The rise of Wikipedia as a free, cooperative alternative to traditional hard-copy encyclopaedias is clearly analogous with the guitar tablature websites. We may also note the narrowing gap between what is possible at a well-resourced university and what can be achieved at resource poor universities, particularly those in developing countries. Researchers who live in lands distant from major research centres can get the very latest discussion papers online rather than waiting months for them to arrive by international surface mail. In addition to trends towards open access to books and journals, discussed earlier in this paper, we might expect a raising of average standards due to the diffusion of computer capabilities and software (econometric work that today can be undertaken on laptop computers in developing countries would have been impossible not so long ago without a 'super computer') and access to PowerPoint presentations, video and audio recordings of performances by leading scholars. In future, where a standard curriculum is being taught in a highly programmed manner, the McDonaldization of higher education[3] may reach the stage where students get their lectures remotely from star performers as part of the package offered by major textbook publishers, with most academics relegated more to giving tutorials or providing one-on-one consultations. For more specialized, advanced-level training, students may be able to go to the websites of key academics directly and hear about their contributions via their own videos. The main barrier to this would be attempts by universities to protect their competitive advantage by barring their staff from making their work accessible to outsiders via their websites. Even with this impediment we might expect their public performances to be captured and loaded on to YouTube in much the same way as happens with live performances of bands or news captured on mobile phones (as with events in Iran following the 2009 elections).

Now let us turn to the downsides of easy access to knowledge. First, it should be noted that easier access to information that can help develop a particular capability will not necessarily result in people who wish to develop the capability being more likely to succeed in doing so once they are armed with copious amounts of this information. The problem is that information does not instantaneously translate into know-how, rather it has to be worked with until new patterns of neural connections are established: eventually, things 'click' and we get 'the knack'. What people achieve when they try to acquire a particular capability is a result of not just their innate talents and access to relevant information but also of their inputs of time, effort and attention. If there is no increase in time for study, a spectacular increase

in access to knowledge resources means that the person seeking to acquire the capability risks a diversion of attention and a general dilution of insight due to dallying with many knowledge resources rather than persistence with just a few. This can easily be seen if we return to the guitar tablature example. With so many role models to choose from at tablature websites, some would-be technically proficient guitarists may explore many achievements of many great players briefly without mastering a single piece of their work: each time obstacles are encountered with one piece, they may try others, until they encounter obstacles there and move on yet again. Each time they do so, the cognitive dissonance[4] between their goal and their failure to get very far with each piece they have so far explored can be resolved by telling themselves, in effect, that since each piece is different, they are moving towards their goal by adding to their repertoires of technical 'chops'. The trouble is, they may be failing to perfect many of them by repeated practice.

It would be much harder to delude oneself about progress if one only had access to a single virtuoso work to try to master. Here, when progress is slow and attention wanders, it is more likely to wander to another page of the music and eventually all the parts may come together. Moreover, the more time a person has committed to achieving mastery of a particular piece of codified knowledge, the more we would expect sunk cost biases[5] to keep them on track to completing the task: the thought that 'I've been trying to play this for weeks now, so I'm not going to let it beat me' will tend to result in the dissonance between aspirations and attainments being removed by a focus on the progress that *has* been achieved, since if there are no alternatives the putative benefits of switching to them cannot be raised.

In the area of academic learning, both by students and by their professors, the information revolution has changed expectations about how well researched a piece of work should be. In the past, even a copyright library might not be adequate to satisfy relatively small classes of students since there might only be a single copy of a work that all wished to read. Now, if journal articles are available online there is no 'I couldn't get hold of it' excuse for students who fail to follow up suggestions to read particular works. Similarly, researchers submitting papers to journals know that the more comprehensive their reference lists are, the bigger the chance they have of keeping referees at bay: with modern search engines and, in the absence of a library copy, the ease of obtaining an electronic version in conference paper or discussion paper form, or directly from the author, there is no excuse for not referring to relevant work.

Information technology has dramatically reshaped the terms of the trade-off that researchers face. As Heiner points out,[6] the boundedly rational chooser risks making errors for two reasons: (i) failing to search far enough for relevant information; and (ii) poor processing of information that has been gathered. The more sources we consult, the less attention we give to the average source, so we reduce the risk of oversight at cost of increasing the risk of failure to gain insight from what we examine. The overabundance of online knowledge may thus present major barriers to effective knowledge acquisition. A chronic feeling that we ought to be consulting more sources is at odds with the reflective thinking that is conducive to the generation of new knowledge. When a search engine has uncovered a vast list of possible sources of relevant knowledge, and access to each new source is only a mouse-click or two away, the temptation is to move on when a work becomes difficult to fathom, possibly taking on board its conclusions without critically examining their foundations. Scholars may thus end up merely surfing the literature,

rather like television viewers with remote control devices and a huge range of pay-TV channels vying for their attention who rarely watch a complete programme.

In the past, students often had to make do with reading from only a few sources and then had to use their developing analytical skills to construct an answer for themselves in which they argued their own case. Now, they essentially seek to *report* what others have said on the matter at hand, adding little value in terms of their own *analysis*. Even short assignments have copious lists of references, most or all of which could not have been examined in any depth, and in the worst cases what is submitted is essentially a collage of article abstracts with linking sentences written in a fumbling style that gives the game away. In other words, what has happened is that the students are largely acquiring information and 'know-that': they can report on who has said what but their failure to invest time in understanding the basis for what has been said results in them adding little to their capacity to 'fly solo' in the disciplinary area and do analysis of their own. After graduating and becoming, say, policy analysts, the reports that they write for their bosses will be largely that—i.e. reports—extrapolating from the findings of others rather than based on the application of principles that they have mastered.

Modern academics likewise operate very differently as scholars from their counterparts born a century or so ago. For example, an academic economist now can potentially draw upon articles in well over 1,000 journals at various levels of quality. Even at the top end, there are far more than they could hope to browse on a regular basis (the Australian Deans' list has 40 ranked as A*, and over 80 ranked at A). By contrast, in the 1930s it was possible to keep up with the literature via regular browsing of the *American Economic Review*, *Economic Journal*, *Journal of Political Economy*, *Quarterly Journal of Economics* and *Review of Economic Statistics*, with the emergence of new journals such as *Economica*, *Oxford Economic Papers* and *Review of Economic Studies* as a glimmer of the flood that would follow. Not only this, but articles in those journals had relatively few references to follow up, usually in footnotes, in contrast to the long lists at the end of today's articles.

In those days, a typical article did not have an abstract and its wordy nature could require a serious investment of time to work out its key messages. A serious scholar, such as George Shackle (1903–92) would return to particular books or articles on many occasions, checking his or her understanding and seeking to glean further insights.[7] The narrow range of works on any topic was conducive to deep, reflective reading. Because there were fewer works to draw upon, creative thinking and the ability to work things out from first principles were much more important capabilities to possess than skills in uncovering and filtering the contributions of others.

The advent of information technology for libraries and its wider application in the Internet may have seemed a blessed relief to academics who, by the 1990s, were having to deal with the proliferation of journals and books. Systems for generating bibliographies by snowballing from the footnotes and reference lists of previously consulted works were by that stage becoming increasingly unworkable: ways were needed to know rapidly how relevant particular sources might be (hence the rise of requirements for abstracts and emergence of specialist abstracting journals, and survey articles) and for ensuring that potentially relevant sources were not being overlooked due to not appearing in the network of citations because other scholars had also previously failed to discover them.

Today, though, search engines threaten to overload us with sources of relevant knowledge. Given the 'experience good' nature of scholarly contributions, the

scholar needs decision rules to help work out where to allocate the most attention. To some extent the citation rate and knowledge of the author's reputation can serve in this role. However, while a consumer can cope with information overload by ignoring most of the information and relying on a familiar brand or a 'best buy' from a trusted source of product reviews, a scholar is supposed to leave no stone unturned and is expected to 'review the literature' as part of the modern academic writing format. Ultimately, however, the academic, like the consumer, has to shop for knowledge on the basis of familiar brands (of authors and/or journals) and by copying others (checking sources in order of citation 'hits'). Hence although publishing in low-circulation journals in principle should nowadays not be the kiss of death for a contribution to knowledge, since one's work is more likely to be discovered and is more easily accessed, it is likely that in practice little has changed in this respect.

Information Technology and the Creation of Codified Knowledge

It is natural for economists to assume that the advent of word processing technology, coupled with referencing software such as Endnote, has greatly increased the productivity of those who produce knowledge. In the early 1980s my first books were written by hand and then typed by a secretary who also typed my reference lists from stacks of index cards. I had to wait a few days for each finished chapter or paper and if I wanted to make any changes these required cunning use of white-out fluids and cutting and pasting of a physical kind, aided by 'magic tape'. Modern technology liberates us from all that and seems to eliminate expenditure on typists and stationery. Producers of knowledge may not be able to type as fast or as accurately as specialist typists but disintermediation does not matter if they can type as fast as they can generate their ideas and can check their mistakes as they go, rather than having a separate proof-reading stage after work comes back from the typist.

This view ignores the impact that being liberated from the old technology has on the process of writing. The freedom to make changes without imposing major re-typing or physical cutting and pasting costs changes how we write and may result in much more time being spent on writing than in the days when we actually wrote on paper or used a traditional typewriter. Academics who repeatedly made changes to their manuscripts used to have not merely to wait for successive versions of their work to be typed but also faced the prospect of 'frown costs' being imposed by their typists whom it was wise to keep onside in order to get urgent typing done at short notice. When changes were expensive to make, it paid to invest in getting a very clear idea of what one wanted to say before attempting to write it down, and writing tended to be a much more linear process than it is when we have the luxury of word processing. To be sure, it was always possible to write small portions at a time, out of sequence, and then assemble them into a single document in the right sequence before handing it over to a typist. Likewise it was always possible to cross out passages and rework them, or physically cut and paste larger sections. However, because these costs of changing and polishing our manuscripts were greater than if we were to word process exactly the same set of words, and because messy typescripts were more likely to cause confusion for our typists, the incentive was much greater to get it right first time. In other words, prior to word processing, the process of codifying knowledge was much more akin to recording a live performance by actors or musicians, whereas now it is much more like making a movie or

recording music in a studio where there can be many takes and much of what is produced is thrown away in the process of editing. In yet other words, we might say that writing used to be akin to downloading from our minds to paper, whereas with word processing it has become more like gradually sculpting a statue out of a block of stone or layering on more and more paint to create an oil painting.

Seen thus, it is by no means obvious that the world of word-processed knowledge has the kinds of efficiency advantages that it first seemed to offer. There is greater potential for avoidable cost escalations as writers seek to create the optimal product, rather than merely aiming for something that is good enough to get their message across. Academics may look like they are being productive as they beaver away in front of their computer screens but perhaps they would produce more if they spent more time ruminating on walks around their campuses or college gardens in the manner of their intellectual ancestors. Instead of first getting their ideas clear in their heads by careful thought over a long period, aided by discussions with colleagues, the modern academic may often start writing with only a rough idea of what the finished product will look like and then composes and clarifies it in the process of writing and editing. The creative freedom that word processing provides may lead to more time being spent developing, on-screen, ideas of less significance than would be the case if we felt a bigger pressure to engage in due diligence with our ideas and tried them out on colleagues in lunchtime workshops. The modern academic is often far too busy to save time and improve quality by these old fashioned means and instead will probably be taking lunch in front of a computer screen whilst wrestling with a paper whose writing began prematurely.[8]

The Durability of Knowledge

McNeely and Wolverton[9] provide many examples of how the choice of storage medium can affect the durability and accurate dissemination of knowledge. Carvings in stone from which rubbings can be made serve far better than writing on bamboo strips that decay more rapidly and have to be copied manually. Electronically stored knowledge can be copied without error, but it may be far less durable than knowledge that is carved in stone. Hyperlinks that no longer work, coupled with the frequent updating of websites, have made electronic knowledge difficult to rely upon and pin down: when we refer to it, we conventionally signify its instability by noting not merely its web address but also the date upon which we accessed it, something that we do not need to do with formally published knowledge. Multiple editions of books can be delineated by reference to a specific edition, whereas superseded web-pages simply vanish forever. I have personal experience of this phenomenon. With Bruce Littleboy, I produced the proceedings of the 2007 Conference of the History of Economic Thought Society of Australia as a very professionally presented e-book, with only a very small print run of the physical version.[10] We intended to set a precedent for what could be done and presumed that it would be kept on the HETSA website indefinitely, like the HETSA journal *History of Economics Review*. But no: when the 'conferences' part of the website was updated to provide details of the latest conference, the link to this volume was taken down, with no explanation. Fortunately, some of the hard copy versions ended up in libraries.

Another kind of vulnerability that can afflict electronically stored knowledge is due to advances in information technology that lead to changing standards for

both hardware and software (including changes in programming languages). Just as film archivists may have to go to extraordinary lengths to capture old movie footage from decaying acetate film, so tapes of data generated from the moon landings 40 years ago are difficult to revisit today due to the physical decay of the tapes themselves and the difficulty in finding hardware that will read them. While books and papers may suffer gradual physical decay, they may outlive electronic forms of information storage and be more easily copied when they are several decades old. Although data stored on computer disks in the 1980s may not be prone to the kind of physical decay that befalls data stored on tape, it may be impossible to access nowadays unless one can find an old enough computer (for example, one that actually has a disk drive of the right kind and uses the right operating system and/or programmes).

In the light of this, we perhaps should not be confident that the pdf standard we take for granted today is one with which all future systems will be backward-compatible. At best, we are likely to find that many electronic documents from the period 1985–2000 have been lost due to system incompatibility and the junking of computers that still worked but had fallen behind in technological terms, just as many classic television programmes are no longer available because their tapes were wiped to economize on then-high storage costs at a time when their future market value was not recognized. Even if Word and pdf documents continue to be readable in decades to come, e-mail exchanges containing significant clues about how ideas developed are likely to prove far more ephemeral than word-processed letters from the 1980s. Thus, in 50 years' time, it may well be far easier for, say, a historian of economic thought to study the work and correspondence of an economist from the turn of the twentieth century than one from the turn of the twenty-first century.

These issues will be compounded to the extent that knowledge is increasingly stored not in libraries but at the personal websites of those who generate it, from where users view their blog entries and download more formal documents. When the person in question dies, the website may die too. This is less likely for the stellar performers—for example, Carnegie Mellon University maintains Herbert Simon's final 2001 web materials at http://www.psy.cmu.edu/psy/fAculty/hSimon/hsimon.html. However, the rest of us will have to hope that an intellectual fan will be enthusiastic enough to get access to our websites and keep our e-contributions alive in much the same way as happens with websites dedicated to the memory of deceased musicians whose record companies do not find it profitable to maintain them. If there is not an unbroken succession of such enthusiasts, wider revivals of interest in a particular scholar's work several generations later will be hindered.

Concluding Comments

Useful lessons for where the information revolution may eventually take processes of creating and transmitting knowledge may be inferred from the way that iPods and YouTube have helped transform audiovisual entertainment. The zero-marginal-cost economics of online academic publishing have a good chance of leading to a world in which scholarly knowledge is as freely available as MP3 files appear to be for teenagers. But, as with listening to music, there is still a cost to reading even if we do not have to pay for content, namely, the opportunity cost of our attention. We can attend to existing knowledge deeply or widely, or we can attend to creating new knowledge from what we already know, but each comes at the cost of foregoing the others. The impact of the information revolution on how

we allocate our attention needs to be carefully studied before conclusions are drawn about its implications for the future of knowledge.

Notes and References

1. For a discussion of the latter, see p. 289 of Peter E. Earl, 'Principal–agent problems and structural change in the advertising industry', *Prometheus*, 9, 2, December 1991, pp. 274–95.
2. The distinction between 'know-that' and 'know-how' comes from Gilbert Ryle, *The Concept of Mind*, Hutchinson, London, 1949.
3. See George Ritzer, *The McDonaldization of Society*, Pine Forge Press, Thousand Oaks, CA, 2000.
4. Cognitive dissonance theory was introduced by Leon Festinger, *A Theory of Cognitive Dissonance*, Stanford University Press, Stanford, CA, 1957. For a summary of it and its economic applications, see Peter E. Earl and Robert A. Wicklund, 'Cognitive dissonance', in Peter E. Earl and Simon Kemp (eds), *The Elgar Companion to Consumer Research and Economic Psychology*, Edward Elgar, Cheltenham, 1999, pp. 81–8.
5. See pp. 47–50 of Richard Thaler, 'Toward a positive theory of consumer choice', *Journal of Economic Behavior and Organization*, 1, 1, 1980, pp. 39–60.
6. Ronald A. Heiner, 'The economics of information when decisions are imperfect', in A. J. MacFadyen and H. W. MacFadyen (eds), *Economic Psychology: Intersections in Theory and Application*, North-Holland, Amsterdam, 1986, pp. 295–350.
7. Shackle's study habits and his ways of annotating what he read are discussed by Stephen C. Littlechild, 'The Shackle papers at Cambridge', in Peter E. Earl and Stephen H. Frowen (eds), *Economics as an Art of Thought: Essays in Memory of G.L.S. Shackle*, Routledge, London, 2000, pp. 334–9.
8. For a useful review of some related research findings regarding the impact of the Internet on brain functioning, see Nicholas Carr, 'Is Google making us stupid?', *The Atlantic Online*, July/August 2008. Available at: http://www.theatlantic.com/doc/200807/google.
9. Ian McNeely and Lisa Wolverton, *Reinventing Knowledge-From Alexandria to the Internet*, W.W. Norton, New York, 2008.
10. Peter E. Earl and Bruce Littleboy (eds), *Regarding the Past: Proceedings of the 20th Conference of the History of Economic Thought Society of Australia*, 11–13 July 2007, School of Economics, University of Queensland, St Lucia, QLD, 2008.