4

The Bit and the Pendulum: Balancing the Interests of Stakeholders in Digital Publishing

Information doesn't want to be free.
It wants to be paid for.

Member of the audience, Computers, Privacy, and Freedom Conference, March 1997

The drive toward digital publishing reflects our need to be heard. It speaks powerfully to the dream that everyone ought to have instant access to the best ideas, the most creative works, and the most useful information. On a global network publishers can distribute digital works nearly instantaneously at low production costs, giving consumers the convenience of twenty-four-hour automated shopping.

Technology does not, however, exist in a vacuum. Even if all the technological obstacles to trusted systems described in chapter 3 were removed, serious social and legal issues related to digital publishing would remain. At present, then, the potential for digital publishing remains just that—a potential. The market remains nascent because the medium has failed, so far, to balance the interests of important stakeholders. In this chapter, therefore, we consider the dream of digital publishing and the co-evolution of technological, business, and legal innovations needed to balance those interests.

The Pendulum Swings

Computers and the digital medium itself are sometimes seen as the major barriers to digital publishing. When personal computers and desktop publishing first appeared in the early 1980s, many publishers saw digital publishing as too risky. At the time, numerous factors, such as the lack of an installed base of computers and the high costs of production, reinforced

publishers' doubts. Another deterrent was fear of widespread unauthorized copying. Realistically concerned about losing control over their intellectual assets, many publishers completely avoided the digital medium. From their perspective, the pendulum representing the balance of power between creators and consumers had swung too far toward consumers.

In the late 1990s, several vendors—including Folio, IBM, InterTrust, Xerox, and Wave Systems—introduced trusted systems for digital publishing. These systems vary in their hardware and software-security arrangements, but they all automatically enforce the specific terms and conditions under which a digital work can be used. For example, rights may be time-limited, or offered to different people and groups (e.g., members of affiliated book clubs) at different fees. Some trusted systems differentiate among such uses as making a digital copy, rendering a work on a screen, printing it on a color printer, or extracting a portion of the work for inclusion in a new work. When asked to perform an operation not licensed by a work's specific terms and conditions, a trusted system refuses to carry it out. So dramatically are trusted systems expected to alter the balance of power between publishers and consumers that some observers have suggested that the pendulum of power is now swinging too far toward publishers.

Copyright and Paper Publishing

Copyright law and user practices derive from several centuries of experience with publishing on paper. In order to promote the creation and distribution of useful works, copyright law grants rights holders—authors and publishers—certain exclusive rights. Authors' exclusive right to control the reproduction and distribution of their works protects their interests and those of publishers, who invest heavily in the development of works as well as in their printing, warehousing, and distribution. During the term of the copyright, the law forbids a second publisher from undermining the market by selling competing and unauthorized copies. Copyright law also addresses the public interest in the free flow of goods and information; for example, through the first-sale doctrine—which ensures buyers' right to dispose of a paper book in any way they wish—and the convention of fair use—which permits the quotation of limited portions of a book for review or scholarly purposes.

Copyright law, by itself, does not prevent unauthorized copying. Where the enforcement of copyright law is imperfect, however, technology limitations and the economics of paper publishing help to check infringements. Although photocopying makes it easy for an individual to make a single copy of a work for personal use, it generally does not facilitate large-scale copying and distribution. The offset presses and other mass-production equipment used by printers and publishers generate books and magazines that are less expensive and higher in quality than photocopied materials. Furthermore, the costs of storing and distributing thousands of paper copies are usually prohibitive for individuals not in the publishing business. Thus publishers of printed materials are protected from infringement by well-funded rogue publishers by legal remedies, and from large-scale photocopying by individuals by quality considerations and the relative economies of scale. Together, these forces create a point of balance between publishers and consumers of paper-based materials.

Copyright and Personal Computers

This balance of power in the paper medium does not, however, directly translate to the digital medium. Even though digital works can be expensive to develop, copying them is essentially free. A consumer with a personal computer and a laser printer can produce a digital copy of a work as inexpensive and as high in quality as the publishers' original. Furthermore, during the PC revolution, community support for copyright has been crucially different for paper and digital media. Whereas copyright in paper media is strong and well established, support for copyright of digital works is weak, or even absent.

Many key institutions have long been active in advocating copyrights for work on paper. In the United States, these include the Library of Congress, the Copyright Clearance Center, and the Association of American Publishers (AAP). Worldwide, although there are national differences in copyright law, an international body (WIPO, the World Intellectual Property Organization) is active in promoting treaties and standards to harmonize national laws. Overall, copyright has worked well enough to support the established book, magazine, and newspaper publishing industries.

As paper publishers began to consider digital publishing, however, they faced resistance from the computer community. Many people in that community believe deeply that computer software and information ought to be free. Although this attitude is contrary to ordinary business practice and market theory, it has grown naturally out of the history and environment of computer science.

Before the rise of the personal computer in the 1980s, digital publishing was mostly the province of academics and scientists, who shared computers and used them to distribute the results of their scientific work. Academic publishing is a special case, in that the sharing of research results is rooted in a philosophy favoring the free exchange of ideas and the open search for knowledge. Academic journals are also unusual in that authors are not paid for their articles, as they are in commercial publishing. Authors, for whom scholarly citation is important for academic reputation and career advancement, often pay to be published through page charges, rather than being paid for their writings.

Software developed for various scientific purposes in the computer-using academic community was created by and freely exchanged among many groups of users. Widely used software enhanced its creator's reputation and standing in the computer community. This was a natural extension for a community in which sharing results and peer review are key to academic advancement. The free exchange of programs that began in the 1960s accelerated the development of computer science, the field directly concerned with the creation and study of the algorithms used to construct software programs.

The wider PC community inherited from the academic community the values supporting free exchange of information. When the PC community first emerged as a hobbyist fringe of academic computer science, these values served it well. Many hobbyists built their own computers and traded programs. However, as the personal computer culture—and industry—matured, the underlying assumptions supporting free sharing no longer applied. Fewer users created their own software or built on each other's programs; instead, they bought software from publishers. Publishers did not rely on academic grants for support; they had to make a living selling their software.

To keep prices down, software publishers amortized the costs of development and production over their user base. They devised several measures—such as rigging computer disks in various ways—to make it difficult

to copy programs. However, even legitimate customers found such protection approaches too inconvenient, and they were eventually dropped. Ultimately, large software publishers realized that copying could be good for business. People often learned a software program by using pirated copies, then later bought legitimate copies of newer versions. This process created a market dynamic that helped large publishers dominate the field and marginalize smaller ones.

By the early 1980s software publishers were operating in a legal regime that provided, at best, uncertain safeguards against the copying of digital works. The courts pondered the extent to which copyright ought to protect the structure, sequence, and organization of computer programs (e.g., in *Computer Associates v. Altai 1992*). Certain well-known computer copyright cases—such as Apple's dispute with Microsoft over the Macintosh graphical user interface *(Apple v. Microsoft 1994)*—took many years to resolve. Meanwhile in other electronic venues such as videocassette recording, court challenges suggested that certain kinds of copying were permissible on the grounds of fair use *(Sony v. Universal Studios 1984)*. And, as noted in chapter 3, Congress refused to pass laws to regulate copying of satellite television programs until broadcasters developed adequate anticopying devices.

Comparing computers to leaky bottles, John Perry Barlow (1994) argued that once a digital work is created, it will inevitably be copied.

Copyright and Trusted Systems

Beginning in the 1990s, however, computer scientists realized that computers could become part of the solution to the copyright problem they had reputedly caused. The key was development of trusted systems technology.

The two main ideas behind trusted systems are (1) that the terms and conditions governing the authorized use of a digital work can be expressed in a computer-interpretable language; and (2) that computers and software can be designed to enforce those terms and conditions. Xerox's DPRL (Digital Property Rights Language) is an example of such a rights language.

Digital rights cluster into several categories. Transport rights include the right to copy, transfer, or loan a work; render rights pertain to playing and printing a work; and derivative-work rights govern excerpting portions of a work, limited editing of it, and embedding parts of it in other works. Other

rights govern the making and restoring of backup copies. With trusted systems, a publisher can assign these rights to a digital work and stipulate the fees and access conditions governing the exercise of each specific right.

Trusted systems enforce the assigned terms and conditions and allow exchange of the work only with systems that can prove themselves to be trusted systems through challenge-response protocols. Trusted systems thus form a closed network of computers that excludes non-trusted systems and that collectively supports use of digital works under established rules of commerce. When digital works are sent between trusted systems, the works are encrypted. When they are rendered by being printed on paper, displayed on monitors, or played on speakers, the system can embed in the signal machine-readable watermark data that make it easier to trace the source of any unauthorized copies.

In general, the higher the security of a trusted system, the higher its cost. High-security trusted systems can detect physical tampering, set off alarms, and erase secret key information. Intermediate-security trusted systems have more modest physical, encryption, and programmatic defenses. Using challenge-response protocols, all trusted systems can recognize other trusted systems and determine their security levels. This enables publishers to specify for each work the security level of the trusted systems allowed to receive it. A sensitive industry report, for example, might require an expensive and secure corporate trusted system with advanced security features. On the other hand, a widely distributed digital newspaper subsidized by advertisements might require only a modest level of security easily attained by home computers.

Trusted Systems and the Balance of Interests

There are many stakeholders in digital publishing. Besides the federal government, U.S. copyright law focuses on two parties or categories of people: rights holders (that is, the authors and publishers who hold the copyrights) and the public. Even in paper-based publishing, however, there are multiple intermediaries, including wholesalers, bookstores, used book stores, and libraries. Trusted systems, which delegate enforcement and control to computers, introduce other third parties: trusted system vendors, financial clearinghouses, and national governments. Thus the use of trusted systems complicates the balance of interests by introducing new stakeholders.

Further, the use of trusted systems to enforce terms and conditions provides a much finer grain of control than copyright law, and it moves the legal basis of protection toward that of contracts and licenses. This more precise degree of control distinguishes among usage rights for copying, loaning, printing, displaying, backing up a work, and so on. It also provides for the identification of specific users, rendering devices, and usage fees. In addition, trusted systems bring about a finer grain of control by enabling rights holders to monitor transactions for the usage of works.

In the following sections we consider the sometimes competing interests of stakeholders and the technological and institutional implications of those interests in digital publishing that uses trusted systems. In particular, we contrast copyright-based and contract-based protection of digital works and examine how issues such as fair use, liability, and national borders are likely to play out in a trusted systems regime.

Copyright Law

As excellent summaries of U.S. copyright law are available elsewhere, we will not recapitulate them here. However, understanding certain key aspects of copyright law will clarify how stakeholder interests will influence the design of trusted systems.

Like other forms of legal protection for intellectual property, including patent and trademark law, copyright encourages the creation of intellectual property by granting certain exclusive rights to its creators. Article 1, section 8 of the U.S. Constitution authorizes Congress to create legislation "to promote the Progress of Science and useful Arts, by securing for limited Times to Authors . . . the exclusive Right to their respective Writings." Over the years, Congress has enacted a series of copyright laws, beginning with the Copyright Act of May 31, 1790. The most recent major overhauls of the copyright law are the Copyright Act of 1909 and the Copyright Act of 1976.

Under copyright law, the federal government grants a copyright owner certain exclusive rights in his or her work. These include the right to reproduce copies of the work, to prepare derivative works based on the work, and to distribute copies to the public—by sale or other transfer of ownership or by rental, lease, or lending. For certain kinds of works, such as literary, musical, and dramatic works, exclusive rights to publicly perform and display the work are granted as well.

These rights last for relatively long periods of time. For works created after January 1, 1978, copyright expires fifty years after the death of the author or composer or (in the case of works made for hire) the earlier of seventy-five years from the date of publication or a hundred years from the date of creation. (Legislation passed in Congress in 1999 extends the copyright term to seventy years after the author's death or, for works made for hire, the earlier of ninety-five years from publication or a hundred and twenty years from the date of creation.) Although the term of copyright is long, it is finite; once it is over, the copyright owner's exclusive rights end. Thereafter, the work falls into the public domain, where anyone may publish it freely.

Copyright law attempts to strike a balance among the competing interests of various stakeholders, especially those of rights holders and the public. It addresses this balance by limiting the duration of exclusive rights. It also limits the scope of the legal protection afforded. The Supreme Court has held, for example, that copyright protection does not extend to the "sweat of the brow" invested by a work's creator but only to the author's clearly original contributions. Thus an alphabetically arranged white pages telephone directory may lack sufficient original content to be protected by copyright, even if compiling it required considerable effort (*Feist v. Rural Telephone 1991;* Samuelson 1996b). Further, copyright protection extends to the expression of ideas but not to the ideas themselves. Thus, in the area of computer software, the courts have held that although the code of a computer program can be (narrowly) protected by copyright, its functionality can be protected, if at all, only by patents or laws governing trade secrets, not by copyright *(Sega v. Accolade 1993; Atari v. Nintendo 1992).*

Further, in striking the balance between rights holders and the public, copyright law provides that reprinting portions of a copyrighted work for purposes such as criticism, comment, new reporting, teaching, scholarship, or research is a fair use and not an infringement of copyright. (We consider fair use as it relates specifically to trusted systems in a later section.) Additionally, provisions of the law (sections 108 to 120 of the Copyright Act of 1976) establish a framework—or, some might say, a patchwork—of more specific rights limitations, scope restrictions, and licensing arrangements. Many of these provisions are designed to address the concerns of particular interest groups, such as religious organizations, small businesses,

lending libraries, blind and handicapped persons, cable television stations, and noncommercial broadcasters.

The history of copyright law has been, in part, a history of revisions intended to keep pace with changes in technology and media. In 1790 the copyright law governed only books and navigational charts; today it pertains not only to texts on paper but to all original works in a fixed, tangible medium of expression (e.g., motion pictures, architectural plans, sculptures, and sound recordings). Nevertheless, major revisions of the copyright law have been few and far between; almost seventy years elapsed between the 1909 and 1976 Acts. Even less-fundamental changes achieved through statutory amendment, regulation, or case law often take years and so lag far behind the latest technological developments.

This lag is especially apparent in the era of digital media, which change at an unprecedented rate. Moreover, digital media blur the boundaries between traditionally separate categories of work; thus statutory provisions intended to protect a particular type of work, or to serve a particular set of interest groups, can rapidly and unexpectedly become applicable elsewhere. For instance, is a web page that includes an animation and an audio stream an audiovisual work or a computer program? Or both? Or something entirely new?

Or consider a packet-switched computer network such as the Internet, in which the content-bearing information packets can travel either on cables or over the airwaves. The choice of a transmission path depends on moment-to-moment routing conditions that are unknown, even unknowable, to users. Should provisions of the copyright law pertaining to cable television or to broadcasting apply to the packets? Does it make sense to define a single concept of transmission for the digital media? And, if the answer is less than clear, should the law provide for a new exclusive right associated with all kinds of information transmission, as some have proposed? Yet another gray area involves whether the bits in an information packet traveling across the Internet are sufficiently "fixed in a tangible medium of expression" to constitute a work subject to copyright.

Digital media can also create confusion about who is the creator or owner of a work. Digital audio processing, for example, allows a composer to take digital samples from other (possibly copyrighted) works, process them in various ways, and include them in other, derived works. The source or sources of the sampled works may or may not be recognizable in their

processed form; the samples can be very short, perhaps only a single drumbeat or a single note of music. In the absence of trusted systems, it is difficult, or impossible, to determine the original source of such a short sample. Even if the source can be identified as someone else's copyrighted work, it is not always clear (1) whether the person doing the sampling should pay the owner a royalty, and (2) what circumstances govern whether such borrowing of short segments constitutes fair use. Such digital sampling—or, in the absence of trusted systems, the lack of a precise method for controlling digital sampling—has already posed problems for the recording industry and recording artists. These issues have become even more complicated now that samples can be traded among thousands of musicians worldwide via the Internet.

Without trusted systems, effective enforcement of copyright in the digital media is nearly impossible. Like the effort to plug the proverbial sieve with its thousands of little holes, finding all the little infringement leaks of isolated individuals making copies is too hard and too expensive. Moreover, living with the leaks has its own deep risks. By publishing without copyright enforcement in a community that routinely makes unauthorized copies, rights holders accept the risk that, over time, such copying could become an established practice and could even be sanctioned by the courts as fair use.

In summary, the move toward digital media poses numerous challenges for copyright law and creates great uncertainty for rights holders, especially for would-be digital publishers. That uncertainty has hampered the adoption of the new media and discouraged most publishers from entering them. The impracticalities of enforcing copyright on untrusted, networked systems; the gray areas of legal interpretation for digital works; the lack of fine-grained control in copyright law; and the risk of an emerging legal claim of fair use for digital copying—all motivate aspiring authors and publishers in the digital media to find means other than copyright law for protecting their interests.

Digital Contracts

To consider one such alternative means, imagine a representative scenario of digital publication on a trusted system. The author begins by creating a

work. When it is complete, the author finds a publisher to further develop the work and sell it (or perhaps decides to self-publish). The publishers develop a set of terms and conditions governing use of the work and, using a rights management language like DPRL, specify the time period over which the rights apply. They also determine what rights are applicable to the work: for example, whether printing is allowed, whether the work can be loaned out for free, whether members of a particular book club will receive a special discount, and so on. They may assign different fees for different rights: for example, deciding either to disallow creation of derivative works or to encourage creation of such works as a source of further revenue. Or they may require readers of the work to present proof—in the form of a digital certificate—that they are over eighteen. In DPRL, the statement of each right specifies the type of right, the time period in which the right is valid, the special licenses (if any) required to exercise the right, and the fee. In a trusted system, these rights are associated with the digital work either by bundling them together in an encrypted file or by assigning the work a unique digital identifier and registering the work and its rights in an on-line database.

Owning versus "Renting" Software

Having a specific and detailed agreement about the terms and conditions of a work's use is clearly an advantage to authors and publishers. But why would a consumer prefer such a system to the alternative: paying a single fee to purchase a copyrighted work outright and then using it in accordance with a set of legal standards applying to all digital works?

I would argue that specialized rules like those described above have potential economic advantages for consumers as well as for publishers. In the present software market, a customer typically purchases a software license for a fixed fee; someone who expects to make little use of the program pays the same fee as a person who will use it for many hours a day. In some markets, this situation is bad for both publishers and consumers, because low-usage consumers may decide not to purchase the software at all. In a trusted system with differential pricing and metered use, the amount customers pay for software would depend on how much they use it. Metered use would therefore allow consumers to "rent" the software under

terms flexible enough to provide for decreasing unit costs for increased usage and conversion to purchase if their usage is great enough.

Another economic advantages might accrue to publishers and consumers of digital works through a variation on the first-sale doctrine. When consumers buy a paper book, they receive and own the copy of the book. After they have read it they are free to give it to a friend or to sell it. The first-sale doctrine in copyright law guarantees these rights. In the DPRL language, the analogous usage right is called a transfer right. Customers who purchase transfer rights can send the work from their own trusted system to a second trusted system; when they do so the copy on the first system is deleted or deactivated so that it can no longer be used. Like handing a book to a friend, a transfer operation maintains the same number of usable copies in circulation. The terms and conditions that allow a work to be transferred at no charge are thus analogous to those pertaining to paper books under the first-sale doctrine.

The right to transfer the work without a fee is exactly what a consumer intending to share it serially with others might want. On the other hand, from a publisher's perspective, a free transfer right is a threat to future sales. If all the people who read a digital work need to buy their own copy, the publisher will sell more copies. One solution would be to offer two different combinations of rights with a given work. In one combination, the consumer would pay the standard price and be able to transfer the work without fee—as in the first-sale doctrine. In another combination, the consumer could get a nontransferable copy of the work at a discount price, or later pay an additional fee to transfer it. Consumers who buy the work for their personal use and do not anticipate giving it away after using it might prefer the discounted purchase.

It can be argued that the first sale-doctrine is grounded in an experience with paper-based works that treats them as physical objects independent of their creative content. Like tools, household objects, or automobiles, such physical objects can be resold at the owner's convenience. Enforcing a law to prevent the resale or giving of paper books would be difficult in any case; so the first-sale doctrine makes practical sense. For digital works exchanged on trusted systems, however, these considerations are less relevant. The publisher and the consumer are free to enter into whatever agreement they see as economically advantageous.

Contract Law and Digital Contracts

In many ways, the terms and conditions specified in DPRL are similar to a contract or license agreement for using a digital work. However, although for convenience we refer to such a set of terms and conditions as a digital contract, it differs from an ordinary contract in crucial ways. Notably, in an ordinary contract, compliance is not automatic; it is the responsibility of the agreeing parties. There may be provisions for monitoring and checking on compliance, but the actual responsibility for acting in accordance with the terms falls on the parties. In addition, enforcement of the contract is ultimately the province of the courts. In contrast, with trusted systems, a substantial part of the enforcement of a digital contract is carried out by the trusted system. In the short term at least, the consumer does not have the option of disregarding a digital contract by, for example, making unauthorized copies of a work. A trusted system refuses to exercise a right that is not sanctioned by the digital contract. Over the longer term, consumers or consumer advocacy groups may negotiate with publishers to obtain different terms and conditions; but even then, the new digital contract will be subject to automatic enforcement by trusted systems.

Contract law is a complex subject that we cannot summarize here in any detail. Even so, we need to point out a few of the basic provisions of contract law as they relate to stakeholder interests and the design of trusted systems. First, however, we should note that at the time of this writing, there is an ongoing, controversial effort to add new provisions concerning "licenses of information and software contracts" to the Uniform Commercial Code, the body of statutory law that governs commercial contracts in most U.S. states. If the proposed provisions are adopted into law, they could have significant implications for digital publishing in general and for trusted systems in particular.

Contracts are agreements entered into by two or more parties. In a typical case, the parties negotiate and come to an agreement on the terms and conditions under which each party provides something of economic value to the other. This bargain for the mutual exchange of value is an important part of what makes an agreement a contract. Typically, a contract includes one or more promises backed by what are legally referred to as "valid considerations." The terms of the contract are usually set out in a written

document, and the parties formalize their agreement with the terms by signing and dating the document. In cases warranting extra care, the parties' signatures may also be witnessed by a registered third party (a notary public), who asks the parties for proof of identity and may even take thumb prints or require other forms of personal identification.

A contract is backed by the force of law. Generally, if one party fails to comply with the agreed-upon terms and conditions, the other party or parties can enforce the contract through the courts. However, there are various circumstances under which the terms and conditions of a contract are not legally enforceable. They are preempted, for example, by the provisions of the U.S. Constitution and of certain statutory laws. The copyright law contains a preemption provision (section 301 of the 1976 Copyright Act) that may, in some cases, render a contract unenforceable. The courts have also developed various doctrines that hold certain contractual provisions to be unfair or improper (e.g., on the grounds of fraud, illegality, breaches of public policy, etc.) and thus unenforceable.

Here is one example of an unenforceable contract. It is not unusual for landlords to offer tenants standard rental agreements consisting of several pages of formal legalese. Suppose that a tenant signs such a document without noticing a clause in small print saying that if he eats mushrooms on Tuesdays he must pay an additional one thousand dollars in rent. The landlord throws a party on the next Tuesday and slyly offers the tenant mushrooms. Since such a clause is outside the normal scope of what belongs in a rental agreement, the courts would very likely refuse to uphold it. The courts thus provide checks and balances in contract law by deciding what contracts to enforce and how to interpret the terms and conditions of those contracts.

Checks and Balances in a Trusted System

With properly designed trusted systems, many of these same checks and balances can be available automatically. Consider, again, the case of the author who has finished a work and gives it to a digital publisher, who assigns it a set of terms and conditions. Like the conventional language used in legal contracts (so-called "boilerplate"), digital boilerplate in the form of templates and default conditions can be used to set up a digital contract.

Suppose, then, that the publisher has included some very unusual terms and conditions in the agreement. When the consumer's trusted system is connected to the publisher's trusted system, it first retrieves the terms and conditions of the digital contract and shows them to the consumer. Before the consumer accepts the digital work, his or her system can use a program (a "contract checker") to check for and highlight any unusual conditions in the contract (e.g., high fees for certain rights, unrealistic expiration dates, or other uncommon requirements). Because rights-management languages like DPRL are formal languages of limited complexity, simple grammar, and predetermined meanings, such a check is a straightforward matter for a computer. (As a somewhat bizarre example, consider a contract provision specifying that the consumer can copy a digital work for free but, surprisingly—and inconveniently—must pay ten dollars to delete it.) Before taking delivery of a work on a trusted system, consumers have the opportunity to agree to the terms and accept delivery of the work or to refuse them and not receive it. If a consumer agrees, his or her trusted system can digitally sign an acceptance form, which can be digitally notarized by a third party (a "digital notary") known to both parties.

The sequence of events in this example illustrates several checks and balances in the process. Both the publisher and the consumer can use computational aids to check the normalcy and appropriateness of the contract. More than a labor-saving or time-saving procedure, this approach is also a compensation for the intangible nature of information inside computers. It increases the confidence of both parties that the terms and conditions used by the trusted systems are reasonable.

One Digital Contract = Several Legal Contracts

It is helpful to think of a digital contract not as one contract but as a combination of several distinct legal contracts. There is the contract for access to the copyrighted work itself, and a second one for delivering the digital data, irrespective of whether the data are or can be copyrighted. If, for example, publishers use a trusted system to provide an uncopyrightable database *(ProCD v. Zeidenberg)* or a telephone white pages directory *(Feist v. Rural Telephone 1991)*, they are entitled to charge for this service, even though the consumer could, in principle, get the uncopyrightable data

elsewhere or put together his or her own database. Similarly, digital publishers, like paper publishers, could charge the consumer for delivering a copy of the complete works of Shakespeare, even though they are in the public domain. Like print publishers, digital publishers make life more convenient for the consumer, who must pay for this convenience. Publishers of an uncopyrightable work or one in the public domain cannot use copyright law to dissuade another publisher from offering consumers the same or a comparable work, as they could if the work were copyrighted. Thus, they must depend on trusted systems and digital contracts to maintain a business position.

Finally, there is a third contract implicit in the digital contract: namely, the agreement entitling the consumer to access the network of trusted systems in the first place. This agreement may be arranged between the consumer and the publisher, or between the consumer and one or more network service providers, who may or may not be affiliated with the publisher.

The idea that a digital contract includes multiple legal contracts provides a coherent rationale for enforcement of digital contracts, even those pertaining to uncopyrightable works. For example, suppose that a digital publisher provides a customer with a work in the public domain, such as the complete works of Shakespeare, under a digital contract that prohibits the copying or further transferal of the work in digital form. One consumer, however, is unhappy about this. He knows that the work is not protected by copyright and, when the bill arrives from the publisher, he refuses to pay it, or he sues to get his money back. In court, the consumer argues that charging for a work no longer protected by copyright violates the copyright principle of providing only limited-term monopolies for authors. Therefore, he claims, the digital contract should be preempted by the Copyright Act and held unenforceable. (The consumer might also argue that, because the publisher accepts the consumer's money while providing in return only a public domain work that ought to be available for free, the agreement fails for lack of consideration; i.e., because nothing of value has been delivered.)

The publisher responds that what is being sold is not the work itself but, rather, the service of delivering it. The publisher says, in effect, "Consumer, by dealing with me, you save time and energy and money over other delivery mechanisms, such as conventional bookstores. If I, as a vendor, want to

provide this service to others, I must be entitled to collect revenue for the service, not just for the work itself. Therefore, it is legitimate for me to prevent you by digital contract from transferring the copy of the work I sold you." We think that the publisher has the better argument here. The consumer can still pay the publisher for the right to print out the contents of the book and can then copy the contents—for example, by hand or by scanning with an untrusted optical scanner. Moreover, other publishers can produce similar books containing identical texts, and a not-for-profit library can make these texts available for free. In short, the publisher has not overstepped the bounds of copyright.

Can the Consumer Negotiate?

Another point of possible concern with a digital contract is the extent to which a user can realistically negotiate the terms of the contract. In court cases challenging the validity of so-called shrink-wrap licenses—software contracts saying that by breaking the plastic seal on a package the purchaser agrees to the terms of an enclosed contract—plaintiffs have argued that such licenses give the publisher a power advantage and leave the user with only a "take it or leave it" choice. Many consumers do not bother to read the license. In principle, with trusted systems, it is possible to open a channel with the publisher to negotiate changes in such terms. It is worth noting, however, that one of the main advantages of digital publishing is the possibility of fully automated, twenty-four-hour shopping convenience. In that setting, renegotiating the terms of purchase for a mass market digital work is as unlikely as expecting to negotiate the price of buying a best-selling paperback at a convenience store in the middle of the night. The consumer would simply have to accept the terms offered or postpone the purchase until a human agent is available to renegotiate them.

Trusted Systems and Fair Use

In addition to specific statutory exceptions to the exclusive rights provided to rights holders, section 107 of the Copyright Act of 1976 sets forth four factors to be considered in determining whether a particular use of a copyrighted work is an infringement or a fair use.

1. The purpose and character of the use, including whether it is commercial in nature or for nonprofit educational purposes;
2. The nature of the copyrighted work;
3. The amount of the work used and its substantiality in relation to the copyrighted work as a whole; and
4. The effect of the use on the potential market for or value of the copyrighted work.

Fair use, itself, is not a public right. Technically, it is a legal defense that can be raised when a copyright owner challenges a particular use of a copyrighted work. In a representative case, the fair-use defense works as follows: A copyright owner publishes a copyrighted work. Without seeking permission, a second party obtains a copy of the work and incorporates portions of it in a new work. The copyright owner objects and takes the second party to court, claiming infringement of copyright. In court, the second party argues that the use made of the work constitutes a fair use and should be permitted. The second party might argue, for example, that the excerpt was used in a legitimate commentary or satire. The court has to consider the facts of the particular case in the light of all four factors. In practice, however, the fourth factor, the undermining effect on the market for the work, is often considered the most important.

One of the concerns some have raised about trusted systems is that they might block consumers' access to works they are entitled to use on a fair-use basis. Because a consumer could not extract a portion of a digital work on a trusted system, he or she would not have the opportunity to create the work that would occasion the fair-use defense. Of course, a trusted system would not prevent a user from manually retyping portions of a copyrighted book or from digitally recording excerpts from a copyrighted audio or video work through a computer's microphone or digital camera. It would however, preclude the operations of cutting and pasting excerpts directly from a rights-protected digital work unless the consumer of the digital work has contracted for the appropriate derivative-work rights. This example shows how far, in trusted systems, the pendulum has swung toward giving more power to authors and publishers.

Arguments about fair use for digital works sometimes tacitly (and incorrectly) assume that publishing risks in the digital medium are similar to

those in the paper medium. However, while it is, as discussed earlier, unlikely that an infringer will make and distribute thousands of paper copies of a work, he or she can copy and mail a thousand digital copies with a single keystroke at no expense whatever. In other words, publishers who granted unrestricted access to each and every anonymous user on the basis of fair use would routinely risk the loss of all their copyrighted assets.

The other side of the coin is that fair use serves the public interest, particularly by helping protect freedom of speech. Parody, academic and social criticism, satire, and other forms of speech that rely on the ability to quote from, paraphrase, and modify portions of others' works are essential ingredients of the mix of speech that characterizes a democratic society. Fair use helps ensure that such borrowing can occur, even when copyright holders are vehemently opposed to the use of even short portions of their work in a critique or a lampoon. Without fair use, such rights holders could effectively quash criticism by preventing critics from publishing. The importance of this public interest aspect of fair use is amplified in the digital medium, where the ease of wide-area communications promotes the ability of everyone in society to engage in such free speech and be heard by all. Ultimately, this inclusiveness benefits society as a whole.

The Microtransactions Approach to Fair Use

Is there a way in the digital medium to balance the risks and benefits of fair use for publishers and consumers? One approach to the question argues that pricing and market forces will, in many instances, render the fair-use issue moot. Currently, fair use is a binary decision: either the use of a work is a fair use or it is not. When the courts rule that a particular use is fair, rights holders must forfeit the income from that use; if the use is declared unfair, the consumer must pay to use it. Thus there are high stakes in fair-use cases. The issue is especially awkward when the cost of being honest greatly exceeds the profit expected from using the work. In contrast, with trusted systems, fees for using copyrighted works in the digital medium may be either large or small and can be collected automatically. When market forces prevail, therefore, disputes that might otherwise arise over fair use will be resolved by setting fees at levels appropriate to specific uses.

Copying Music

A scenario about music illustrates how the market might resolve the issue of fair use. Suppose that a consumer wants to copy a short portion of a song from a friend's CD album. Today, consumers can easily make a tape recording for which the record company and the recording artist are not compensated. Such consumers are seldom found out by the record company; but even if they were they would probably argue that copying only a short part of a song is fair use. The company would very likely disagree. The point is that this is an all-or-nothing situation: either the consumer has to purchase the entire CD just to get a few seconds' worth of audio or the consumer pays nothing (because the copying is deemed a fair use) and the record company incurs substantial losses over thousands of such uses.

With trusted systems, though, a middle ground is possible. The fine-grained control offered by trusted systems makes it possible for consumers to purchase exactly the portion of the audio CD they want, and for the record company to charge and collect a fair price for that audio. In other words, to the extent that fair use is the copyright law's response to market failure, well-designed trusted systems can help correct that failure and eliminate the fair-use issue.

Copying Software

Another example of the microtransactions approach to fair use involves a person who buys a computer game or other software program for home use and expects to share it with other family members. In a household with several computers, he or she would probably make multiple copies of the software, one for each computer. Such copying today usually violates the shrink-wrap license that accompanies software. The family might assert, however, that their copying is fair use, because only one copy of the software is used at a time; in effect, they are simply sharing a single copy of the software.

In a trusted systems regime, would software publishers prevent this practice by requiring each family member to pay a separate fee for every use of the software? Such pricing practices would substantially increase the family's costs and provide it with no apparent benefit. Publishers would not necessarily take this approach. Even with trusted systems, they could keep prices at approximately today's level. They might, for example, sell soft-

ware intended for household use with built-in digital contract provisions designed to make sharing the software among family members easy. They would allow a certain number of copies to be made free but would restrict their use to family members and their home computers. A family member who attempted to provide a "free" copy to someone not in the family, or to use it on a trusted system at work, would be prevented from doing so or required to pay an additional fee.

Printing Hard Copies

Yet another instance of the microtransactions approach to fair use suggests how the common practice of photocopying a single page of a book owned by someone else might change under trusted systems. Today, people typically make such copies without the publishers' knowledge. So long as the copying is fairly minimal, it is popularly considered (rightly or wrongly) to be fair use. With trusted systems, a publisher could discourage such unauthorized copying and, at the same time, benefit consumers. For example, the publisher could charge less for the right to view a page of a digital book on a display screen than for the right to print out a hard copy. This would benefit the reader who just wanted to look at one page rather than the whole book. The charge to view but not to print the page might even be made cheaper than the cost of making a photocopy.

A publisher might also charge less for the right to print the page in a form containing a machine-readable watermark than to print it without the watermark. The watermarked version of the printout would be less prone to unauthorized copying, because trusted digital photocopy systems would detect the watermark and charge for the privilege of copying it. Although consumers could still make copies on older photocopiers, the digital watermark would inhibit at least some unauthorized copying.

The Fair-Use License Approach

A second approach to fair use, one not grounded in faith in market forces or microtransactions, would institute fair-use licenses. In this scenario, Joe, a consumer, applies for a fair-use license the same way he might apply for a driver's or a radio operator's license. To earn the license, he has to study the rules of fair use and pass an examination by an appropriate

organization, which we will call the DPT (for Digital Property Trust). The DPT then certifies Joe's identity and issues him a physical certificate as well as a personalized digital license to use on his trusted system. Under DPT, publishers of digital works would assign each work privileged rights that can be exercised by fair-use licensees. Publishers would also declare an insurance limit based on the expected commercial value of their rights to the digital work. Each time Joe exercises a copy transaction to obtain a copy of the digital work, the transaction fee includes an additional small amount—a share of the insurance premium on the digital work. Because of his fair-use license, Joe can exercise privileged rights to the digital work, but these privileged uses might be monitored, logged, and reported (subject to appropriate legal considerations for his privacy). Suppose, then, that Joe sends thousands of usable copies of the digital work to a mailing list on the Internet, and the publisher takes him to court, claiming damages beyond Joe's ability to pay. The court takes into account both the digital contract and the four factors of fair use. If the court finds in favor of the publisher, the fair-use insurance pays at least part of the damages. If the court finds in favor of Joe, the fair-use insurance pays for some or all of Joe's court costs and attorney fees.

The main point of this example is to illustrate that there are different risks and interests surrounding fair use in the digital media. In DPT, fair use is treated as a licensed privilege analogous to a driver's license, rather than as a legal defense. From a legal perspective, this is a substantial reframing of the fair-use concept that takes into account the greater risks of misappropriation in the digital arena. The example implicitly raises several interesting policy issues.

- Does Joe pay for his own license? If fair use is seen as essentially equivalent to free expression (a right), then fair-use licenses should be free, perhaps subsidized in some way by taxes or the publishing industry. If fair use is a privilege, like a driver's license, the potential licensee should pays a modest price for it.
- What rights does a fair-use license grant? In the example above, the publisher decides what additional rights go with a fair-use license. In an alternative scenario, there is a standard set of rights, possibly defaulting to zero-fee versions of all possible rights. Fair-use insurance covers the financial risks to the publisher if the work is turned loose on the Internet.

- Who pays for the insurance? In the example above, a per-transaction fee to pay the cost of insurance would be levied on all consumers of the digital work; this system would automatically scale the cost of the insurance to the popularity of the work. Alternatively, publishers could pay for the insurance, although this would amount to almost the same thing if a per-transaction fee passes the cost along to consumers. By using a per-transaction fee to collect insurance costs, self-publishers especially could avoid making up-front payment of the premium. In still another alternative, the fair-use licensee would pay the premiums.
- Whom does fair-use insurance protect? In the example above, it is intended to protect the publisher against losses. Should there also be a kind of liability insurance or fair-use bonding for consumers, to guard against claims for damages by publishers who accuse them of unfair use? Should there be a deductible on such insurance?
- Should fair-use actions be monitored? According to one view, such monitoring would violate privacy rights. Another argues that because monitoring would enable rights holders to detect violations of a fair-use license, it would be justified. An intermediate position would log and encrypt actions by the fair-use licensee and make these records available to law enforcement only when there is appropriate reason to believe that the law has been broken.

We have seen that the risks to a publisher in the digital medium for unencumbered fair use are much greater than they are in paper publishing. The first approach we considered to preserve the spirit of fair use in the digital medium relies on market forces and microtransactions to make most uses of a digital work very inexpensive. The second, which holds out the possibility of zero-cost fair use (and perhaps better safeguards the right of free speech), balances the risks and benefits by institutionalizing fair use as a licensed privilege backed by insurance.

A recurring theme in discussions of fair use and trusted systems is the fear that publishers will use trusted systems to take advantage of consumers by unfair pricing and that consumers will be unable to mobilize successfully against this practice. However, trusted systems must serve everyone's interests, or they won't serve anyone's. Publishers and consumers alike will be better served if publishers use trusted systems in a way that recognizes and responds to legitimate consumer expectations—for example, by creating digital contracts that preserve traditional notions of fair use. Publishers can either choose to self-regulate or risk being regulated by outside forces:

the legal system, the marketplace, and public opinion. Publishers who fail to consider consumers' interests may come under attack by free speech and civil rights organizations, consumer advocacy groups, and media commentators; they may find themselves the target of boycotts and a public outcry. Effective regulation may then emerge from legislative action. Moreover, digital publishers who get too greedy will find the competition offering better deals; or market forces may push back the entire industry. Consumers will simply continue to prefer works published in more traditional formats, and the market for digital publishing will remain limited—to everyone's detriment.

Accordingly, publishers who want to promote the growth of the market for digital publications will consider interests other than their own, and will make provisions for fair use. At the same time, as trusted systems make it easier for consumers to respect publishers' copyrights than to risk infringement and rely on the fair-use defense, our very notion of what practices constitute fair use will evolve.

Trusted Systems and Liability for Security Failures

Another issue of concern to potential digital publishers is who will be liable for losses incurred through security failures. Copyright law and the preceding scenarios for using digital works on trusted systems focus mainly on balancing the rights of two parties: rights holders and the public. However, because the trusted system technology has to ensure the security of digital works and enforce digital contracts, the manufacturers and vendors of such systems are central to preventing and dealing with the consequences of security failures.

A fiducial responsibility creates potential liability for trusted system vendors. What happens when the security arrangements for a trusted system fail and, without the action or intent of the user, a copyrighted document is released. Is the platform vendor liable? Or consider the case of an individual who purchases a digital work through a trusted system built by vendor A. That system is later used to transfer the work to a trusted system built by vendor B. Subsequently, the work is transferred to a third trusted system built by vendor C. This system fails in some way and releases the work onto the Net. Are vendors A and B liable because their

trusted systems gave a copy of the work to vendor C's system, which proved untrustworthy?

The following scenario suggests one way in which the risks and liabilities for the failure of trusted systems might be handled. Vendor A builds and sells trusted systems that include hardware and software. Before bringing the system to market, vendor A takes the system to an independent testing organization, the DPT (or Digital Property Trust). The DPT tests the system, gives it a security rating, and issues signed digital certificates to be used by the trusted system in its authentication protocols. These challenge-response protocols and digital certificates make it possible for other trusted systems to determine the identity and security level of vendor A's system. They also make it possible to register all the transactions entered into by the system and, in particular, to keep track of which trusted systems have handled which documents.

Then, consumer Joan, using vendor A's trusted system, buys a copy of a digital work. Before offering the work on the Net, its publisher has assigned various rights to it, declared an insurance limit, and decided on the security level of trusted systems required to receive the work. Joan later transfers the work to a friend's system built by vendor B, and, as in the above scenario, the work is transferred again, to a system built by vendor C, which fails in some way (or, perhaps, is tampered with by an intruder); as a result the work is turned loose on the network. All the trusted system vendors, who are part of an industry coalition designed to deal with such problems, then take measures to isolate and limit the damage, and the publisher's document insurance evaluates and pays for the losses.

The example raises several difficult policy issues. Who is liable if a trusted system improperly releases a document because of a design failure? Who is liable if the security of a trusted system has been undermined by tampering or by a computer virus? Who is liable if a computer with outdated security measures participates in a transaction? Is it reasonable for publishers to take the entire risk, even with informed consent about the nature and limitations of the trusted systems? What prudent and appropriate actions might vendors take when a model of their system is apparently compromised? Should there be mandatory periodic testing and upgrading of trusted systems security? Should insurance rates be higher for works on trusted systems of low security than for those on high-security systems?

Over time, security requirements are likely to increase, and system failures are inevitable. When they occur, stakeholders in digital publishing will need to take prompt, concerted, and coordinated action to contain damage and maintain business as usual. The process of determining the cause of a failure can be complex. If the diagnosis and recovery take place in a highly adversarial atmosphere, trusted system vendors may have difficulty cooperating and sharing information in a way that facilitates expeditious containment of a security problem. Creating an industry insurance pool and establishing standards for cooperation can help publishers and platform vendors spread their risks and cooperate amicably.

In summary, the role a trusted system plays in enforcing usage makes vendors of trusted systems a party to the system for honoring intellectual property rights. The example above illustrates an approach that combines security technology with institutional arrangements. That approach is intended to create a business-as-usual marketplace in digital works in which risks are amortized by insurance, concerted and coordinated action by vendors is prompt, and the compliance and security of trusted systems is determined by an independent organization.

Governments as Stakeholders in Digital Publishing

As pundits have observed, international borders are nothing more than speed bumps on the information superhighway. The rapidity of digital communications and their ability to deliver information goods across national boundaries through trusted systems raises several issues of special interest to governments: import and export controls, national security, and taxation.

With regard to taxation, the automatic billing capabilities of trusted systems will almost certainly attract the interest of taxing authorities unless nations agree to treat the Internet as a free trade zone. In principle, of course, automatic billing of taxes is entirely feasible as long as trusted systems are kept abreast of changes in the tax laws of the trading partners.

More problematic is the issue of delineating boundaries in cyberspace. To a large extent, one computer looks like another in cyberspace, and location is not easy to determine reliably. Using intermediate agents, for example, it is possible to disguise the ultimate destination of a digital work.

Furthermore, with portable computers, the actual physical location can change continuously or frequently.

Consider this scenario. A Frenchman carries a laptop computer into the United States and downloads software from a U.S.-based software publisher. Has he exported it yet? Does he export it when he subsequently carries the laptop through customs? What if an American woman in France with a laptop computer she intends to take back to the United States logs onto the Net and downloads some U.S. software. Has she exported or imported software? Is she subject to French taxation? When she returns to the United States with the computer, has she exported the software twice or not at all? The traditional answers to these questions based on current law may or may not make much sense in cyberspace.

One approach to reframing import and export issues is to register computers in a way analogous to the way we now register ships or establish national embassies in foreign lands. If this approach is adopted, export would occur when someone transfers software from a U.S.-registered computer onto a French-registered computer, regardless of the physical locations of the two computers. Trusted systems could carry digital certificates that would authenticate their nation of registry no matter where they are in the world.

Although copyright law tends to be a national concern (though mediated internationally through the Bern convention), the digital medium intimately connects computers in different nations. The possibility of instituting automatic taxation on electronic commerce through a system of national registry is therefore likely to become a matter of great interest to various national governments.

Reflections

With trusted systems, copyright is potentially alive and well in the digital era, as is the balance of interests that copyright represents.

Trusted systems do not, however, exist in a vacuum. They are complemented by and are complementary to the legal, economic, social, and policy frameworks in which they operate. The Internet edge for digital publishing arises from the interactions among these frameworks and is driven by the desire of publishers and the public to utilize the Internet for

commerce in digital works. The pushback comes from the need to balance the complex of stakeholder interests through copyright and contract law, market forces, and technology development. Another component of the pushback comes from traditionalists in paper-based printing who resist new business models and seek to impose print-based policies on the digital world. In this situation, chaos may arise from the legitimate, and often conflicting, interests of the stakeholders in digital publishing—authors, publishers, consumers, librarians, trusted systems vendors, financial clearinghouses, governments, and the public.

As we have seen, the pendulum is still in motion. Without trusted systems, digital publishing is at risk. If we have the wisdom to understand and institute the right policy choices, trusted systems may give us the leverage we need to guide the pendulum to an appropriate point of equilibrium.