

This printable version was created under a Creative Commons Attribution NonCommercial ShareAlike license (see www.juliecohen.com)

9

The Structural Conditions of Human Flourishing

Within U.S. legal and policy circles, the discourse of information-policy reform has been organized principally around the themes of access to knowledge and network neutrality. Global discourses of information-policy reform are organized around parallel themes of access and connectivity. Each of those themes has contributed powerful insights to our understanding of the principles that should inform information law and policy. Human flourishing requires not only physical well-being but also psychological and social well-being, including the capacity for cultural and political participation. The access-to-knowledge movement reminds us that enjoying the latter goods requires meaningful access to the resources of a common culture. The network-neutrality movement reminds us that in the networked information age, access and architecture are inseparably intertwined, and that power over the technical conditions of access should be closely scrutinized.

This book, however, has demonstrated the need for a more comprehensive, structural understanding of the ways that the information environment can foster, or undermine, capabilities for human flourishing. Some information-policy problems cannot be solved simply by prescribing greater “openness” or more “neutrality.” The everyday behaviors of situated subjects require spaces where they can be enacted, tools with which they can be pursued, and meaningful legal guarantees in which they can claim shelter. In addition, we have seen that play is a vital catalyst of creative practice, subject formation, and material and spatial practice. Those processes do not follow automatic and inevitable trajectories, nor are they equally robust under all conditions. Facilitating the play of everyday practice requires attention not only to information accessibility and network neutrality, but also to the semantic structure of the networked information environment, and more particularly to the interstices within systems of institutional and informational meaning. Both the legal specification of information rights and the design of information architectures should be guided by the need to preserve room for play in the use of cultural resources, in the performance of identity, and in the ongoing adaptation of places and artifacts to everyday needs.

Beginning with the centrality of the play of everyday practice, this chapter derives three subsidiary principles that should inform the design of legal and technical architectures. The first principle remains that of access to knowledge; without the raw materials necessary for social and cultural participation, one cannot participate meaningfully in the development of culture and

community, and without access to the appropriate networks and tools, one cannot partake of the resources that the networked information society has to offer. The second and third principles, however, move beyond access to specify other structural attributes of the networked information environment that are necessary to preserve room for the play of everyday practice.

The second principle, operational transparency, seeks to render the network's geographies of accessibility and inaccessibility less opaque—to counter the trend toward seamless, inscrutable design. Operational transparency entails a set of practices designed to put users themselves in a better position to engage in processes of boundary management and to exercise situated creativity with respect to the network's constituent protocols and processes. To take full advantage of the network's potential to enable human flourishing, network users need meaningful information about how the network and its constituent artifacts and protocols work as well as access to the processes in which network standards are designed.

The final principle concerns the location of the boundaries that define the scope of copyright, privacy, and (un)authorized access to information technologies. To preserve room for play, those boundaries should afford sufficient freedom to access and repurpose cultural and technical materials, and should reserve to individuals and communities sufficient control over both personal information and the experienced boundaries of personal space. This mixture of freedom and control is achieved most effectively when regulatory architectures are characterized by a condition that I will call semantic discontinuity. Semantic discontinuity refers to gaps and inconsistencies within systems of meaning, and to a resulting interstitial complexity that leaves room for the play of everyday practice. In an increasingly networked information society, maintaining those gaps requires interventions designed to counterbalance the forces that seek to close them.

On Enabling Capabilities: The Materiality of Play

Let us begin by returning to the theory of capabilities for human flourishing introduced in Chapter 1. There I asserted that moving beyond the foundational assumptions of liberal individualism would enable more precise specification of exactly what the capabilities approach requires the information environment to provide. It is time to make good on that claim. To enable capabilities for human flourishing, the material and informational infrastructures of the networked information society must afford sufficient room for creative, material, and identity play. Fulfilling this condition requires rules about information access and use that accommodate the materiality and serendipity of everyday practice.

As articulated by its leading advocates, the theory of capabilities for human flourishing begins with a positive conception of human freedom, namely, that human beings cannot attain and enjoy freedom in the truest sense unless a variety of basic needs are met. Nussbaum develops a list of ten such needs or "core capabilities": life; bodily health; bodily integrity; senses, imagination, and thought; emotions; practical reason; affiliation; care for other species and the natural world; play; and political and material control over one's

environment.¹ As one might expect, some items on the list relate to the requirements for physical subsistence. Notably, however, many items address the requirements for moral, emotional, and intellectual subsistence. In particular, many move squarely into the domain with which this book has been concerned: the relationship between the information environment and the networked selves who inhabit it. According to Nussbaum, the ability to use one's senses, imagination, and thought encompasses a right to participate in culture by "experiencing and producing works and events of one's own choice." The ability to play includes the ability "to enjoy recreational activities." The ability to exercise practical reason requires the capacity "to form a conception of the good and to engage in critical reflection about the planning of one's life," and the ability to exercise control over one's environment requires the capacity for effective political participation².

Within information-policy circles, the theory of capabilities for flourishing has become identified with the normative and political claims of the access-to-knowledge (A2K) movement. The importance of capabilities for cultural, moral, and political participation has inspired efforts to develop a more detailed understanding of core informational capabilities and to relate those capabilities to features of the informational and technological environment. Lea Shaver identifies "the ability to access, utilize, and contribute to knowledge" as a distinct capability with its own set of entailments.³ Shaver develops a five-part taxonomy: education for informational literacy, access to the global knowledge commons (including both Internet access and linguistic capability), access to knowledge goods in concrete form, an enabling legal framework (including both laws about intellectual property and laws guaranteeing freedom of expression), and effective innovation systems.

In both Nussbaum's list of core capabilities and more specific efforts to theorize specifications for A2K, however, two important dimensions of the relationship between the information environment and human flourishing remain underdeveloped. The first concerns the materiality of artifacts, architectures, and spaces. We have seen that the human experience of the information environment remains fundamentally embodied and materially mediated. Nussbaum's discussion of control over the material environment presumes that such control is exercised principally through the institution of property, while her conception of practical reason stresses a moral agency that seems only indirectly connected to material practice. Scholarship within the A2K movement has been instrumental in demonstrating that cultural and technical innovation does not invariably require property incentives and that human flourishing requires broad access to the fruits of both cultural and technical innovation. That account of human flourishing, however, still contains an implicit gulf between the intellectual and the material. As we have seen, the intellectual, moral, and material practices of situated individuals and communities are inextricably intertwined. A list of core capabilities for human flourishing therefore should include a non-property-based conception of material agency directed not (only) toward innovation, but also and more fundamentally toward advancing the cultural and moral interests of situated subjects.

The second undertheorized dimension of the relationship between the networked information environment and human flourishing relates to the role of play. As we have seen, most U.S. legal theorists of intellectual property and

privacy, including those affiliated with the capabilities approach, articulate conceptions of play that align with the commitments of liberal theory—that is, accounts of play as a purposive, internal, and unknowable activity. Nussbaum’s description of “being able to use imagination and thought” similarly presupposes seriousness of purpose, while her definition of “play” as the ability to engage in “recreation” seems to contemplate mere frivolity. Without question, the ability to engage in deliberate play with cultural goods, identity practices, and material artifacts is important, but it is not enough. And a dichotomy between purposive play and frivolous recreation would be too simple to encompass all the modes of interaction and experimentation that people pursue. The relationship between human flourishing and play is more complex, and understanding it requires a different and less subject-centered approach.

Chapter 2 developed a broader conception of the play of everyday practice and argued that such play derives its power from its tactical, flexible quality—from its political and phenomenological in-betweenness. So framed, the play of everyday practice performs a vital role along each of the dimensions that this book has explored. Within the cultural environment, the play of everyday practice is what generates creative progress; progress emerges in a gradual, nonlinear fashion as situated users appropriate, imitate, and rework the artifacts and techniques encountered within cultural landscapes. Within the social environment, the play of everyday practice responds to continual encounters with the new and unfamiliar, and so informs the development and gradual evolution of critical subjectivity. Within technical environments, the play of everyday practice adapts and repurposes artifacts and spaces in ways that serve the tactical goals of situated subjects and communities, and this reservation of authority to shape the material conditions of everyday life promotes both innovation and psychological and social well-being.

In each of these domains, the play of everyday practice has structural entailments. The play of everyday practice flourishes in an environment characterized by both resources and opportunity. It flourishes most fully when neither the content of the resources nor the precise nature of the opportunity is fully predictable, and when there is leeway for experimentation. The play of everyday practice exploits imperfect alignment among the sets of overlapping constraints—institutional, discursive, geographic, and material—that characterize experienced reality. It is a function of the size and frequency of the interstices within grids of fixed meaning and permitted action.

To understand why play and rule structure are inextricably related, consider a hypothetical: A major provider of media content announces that henceforth it will embrace fan culture by creating a venue where fans can “play” with the characters from a variety of popular television shows and movies. It establishes a dedicated Web site, stocks the site with preselected video clips, and authorizes fans who register with the site to remix the clips (but only those clips) at will. It announces plans to add more variety—including both more video clips and greater technical capability—at some unspecified time in the future. Advocates of greater legal freedom to remix think this behavior signals a positive shift in corporate culture, but they caution that it is a poor substitute for a less constrained remix culture that allows experimentation with a broader range of images and meanings. The Copyright Office needs to decide whether to issue a rule allowing so-called vidders to circumvent technical protection

measures that regulate access to audiovisual content. It has statutory authority to grant the exception if it concludes that creators of fanvids are “adversely affected . . . in their ability to make noninfringing uses” of the protected works. In other words, it must decide whether fanvidders enjoy sufficient scope for play when they enjoy only the remix privileges that creators choose to allow.⁴

Strategic decisions by corporate intellectual-property owners to establish defined zones of authorized play are becoming more common, and are not confined to the realm of cultural goods. In some industrial contexts, producers of equipment embodying valuable intellectual property have sought to harness the distributed power of user creativity by authorizing innovation within certain clearly delineated parameters. Such projects do not necessarily accommodate the full range of user interests. Meanwhile, some scholars who have begun to study the cultural dynamics of crowd-sourced innovation argue that the cynical use of play tropes in the networked information environment too often conceals the structured exploitation of unremunerated labor.⁵

These examples remind us that “play” is a term capable of multiple, often competing interpretations. Rules for play can be liberating or infantilizing, productive or stifling, depending on whether they are appropriate to the actors and to the activity. For a variety of reasons, it makes good sense to restrict very young children to playgrounds and to subject children of all ages to appropriate rules about the nature, location, and scope of play. Geographic boundaries and scope rules can become more controversial when applied to adult play. Many meaningful activities that we count as playful occur within sets of constraints—consider, for example, soccer or jazz or poetry or cryptography research. But not all constraints should be valued equally. Critically, focusing solely on whether the actor understands her own activities as playful can’t lead us to a meaningful rule for evaluating the quality of the constraints. Play is related to rule structure in a way that is inverse and inherently interstitial; it is a function of the spaces that the constraints leave unoccupied.

Understanding capabilities for human flourishing as inextricably linked to the structural attributes of the networked information environment—to its semantic and physical architectures and particularly to its interstitial structure—points toward three conditions that are necessary to enable their development. Those conditions are access to knowledge, operational transparency, and semantic discontinuity. Each condition can be satisfied by a range of possible actions, but each requires action within that range. Although some features of the current information environment satisfy the conditions for human flourishing, many do not. Many of the legal and technical developments that this book has explored jeopardize precisely those attributes of the information environment that enable material agency and that facilitate the play of everyday practice. For a host of often well-intentioned and seemingly logical reasons, those developments seek to shrink the interstices of regulatory and informational regimes—to reduce the looseness of fit between individual behavior and the institutional and technical structures that operate to constrain and channel it. The remaining sections of this chapter explore each of the three conditions for human flourishing and identify types of initiatives that might reverse the trend.

Access to Knowledge

First, the development of capabilities for human flourishing requires access to information, including networked information resources. The access must be of sufficient quantity and quality to enable participation in cultural and social life for purposes ranging from the political to the “merely” recreational. In addition, the access imperative encompasses any tools and technologies that are or may become necessary to enjoy access. The access imperative is the subject of a growing body of commentary, so my discussion will be brief and will focus on highlighting major areas for reform.

As a preliminary matter, it is important to note that my formulation of the access imperative excludes several important dimensions of informational capability that are simply beyond the scope of this book. First, informational capability has social and institutional dimensions. Social literacy includes the information necessary to function as a member of a particular community. Such information is acquired by a variety of processes, only some of which involve access to intellectual goods. Institutional literacy includes, for example, market literacy—that is, knowledge about bank accounts, access to credit, and whatever else is necessary to survive and thrive in an economy that is increasingly networked and global. Second, informational capability has a political dimension that does not depend only on access to intellectual goods, but also requires legal safeguards for freedom of expression and political participation. Entire literatures are devoted to each of those topics. Here, I will focus more narrowly on access to the informational resources necessary for cultural, social, and political participation.

Both the attainment of basic literacy and the capability for ongoing cultural and political participation are powerfully shaped by the ways that intellectual property laws, and particularly copyright laws, mediate access to intellectual and cultural resources. Copyright scholars have long recognized that a system of proprietary rights in intellectual goods inevitably creates some “deadweight loss.” Simply put, if authors and publishers are entitled to demand payment, users who are either unwilling or unable to pay will lose out. Scholars have justified the deadweight loss by pointing to copyright’s presumed incentive function. If authors and publishers are not entitled to demand payment for their works, they reason, copyright will cease to function as a meaningful incentive for production of some important cultural goods, leaving all of us the poorer. Neither side of that equation requires perfection, however. As we saw in Chapter 6, the incentive side of the equation requires reformulation; copyright is not the primary motivator of creative practice by situated subjects, but instead is more accurately described as serving an economic-organization function. Furnishing sufficient control for copyright to serve that function does not require unconstrained control over either pricing or use, and satisfying the access imperative does not require free access to everything all of the time. Honoring the access imperative within the framework of copyright law requires only enough paid access for the market-based system of cultural intermediation to work as a practical matter and enough free or lower-cost access to satisfy basic conditions of social justice.

Although the system of global copyright expanded its reach dramatically during the twentieth century, one of the twentieth-century copyright sys-

tem's signal virtues was that it left unobstructed a variety of avenues for obtaining free or lower-cost access to copyrighted materials. Resale markets afforded access to many mass-produced works at reduced prices, while public libraries provided free access to many different kinds of works. Both institutions functioned as mechanisms for implicit price discrimination, crudely sorting users based on their willingness to pay while excluding no one based on inability to pay. The inexact sorting mechanisms of analog copyright do not simply represent analog imperfection; they performed an important access-promoting function because they were inexact. In economic terms, the consumer surplus resulting from that system factored importantly into the social benefits that the system of copyright law produced. By widely distributing access to cultural resources, it created a broad and deep foundation for cultural participation, cultural production, and cultural progress. A more perfectly calibrated system of compensated access to cultural resources, which would supplant some functions traditionally performed by resale markets and libraries, would not produce the same benefits.

Consider the current controversy regarding the Google Book Search project. On one reading of the dispute, Google Book Search represents a collision between two models of full-text access to cultural works, the library and the bookstore. If we understand the bookstore as the default and the library as the exception, we might be inclined to think that a principal *raison d'être* of the library is to fill unmonetized gaps in the system of access. In the era of full-text digital searches and print-on-demand availability, we might expect the two models of access to converge to a significant degree as the reach of monetized access expands. If so, we might favor sharply limiting the extent of uncompensated full-text access available to library patrons. But if libraries perform functions that are important and qualitatively different from the functions that bookstores perform—functions related to distributive justice, to cultural preservation, and to the promotion of serendipitous access—then that stance would not make nearly as much sense. If libraries are not simply gap fillers, then we ought to approach the question of full-text access very differently. Rather than asking how libraries should accommodate the interests of publishers, we should ask how the functions of libraries might best be preserved in a digital world.

A different set of problems is presented when market formation is desirable but is stymied by transaction-cost and holdout problems. In such cases, access can be promoted by regimes of automatic licensing. An example of the difference that such licensing can make comes from the realm of music copyright. In 1909, lawmakers concerned to foster widespread, affordable access to musical works instituted a system of automatic licensing for the preparation of sound recordings. That system, which remains in place today, led to the emergence of a robust, market-based system for access to copies of recorded music. By the late twentieth century, however, the political climate in the United States had become much more hostile to the concept of automatic licensing. Legislation creating a digital public-performance right in sound recordings sharply limited the availability of automatic licensing and also created new legal obstacles to the formation of collective licensing organizations. In part because of the reluctance to embrace automatic licensing more fully, the market for access to digitally streamed music has developed haltingly, and innovative start-up models are rightly perceived by all parties as creating large litigation risks. The irony, of course, is that the major record labels that play the role of market ob-

structionists in today's music-copyright dramas owe their privileged legal status in considerable part to the catalogs that automatic licensing has enabled them to amass.⁶ Copyright reforms designed to promote market-based access should make automatic licensing widespread.

A just regime of copyright also must include strategies specifically designed to enable access by users located in the global South. Scholars and A2K advocates have advocated a variety of measures to address persistent global disparities in educational capabilities. Strategies for ameliorating the North-South access gap involve a range of reforms to national and international laws, and are designed to promote access to basic educational materials for free or at below-market rates.⁷ As a result of pressure from established copyright interests, however, discussions of access reform within the conceptual framework of the General Agreement on Tariffs and Trade have placed great emphasis on safeguarding private bargaining processes and have generally ignored the question whether some markets may require baseline levels of information flow in order to form at all.

The copyright industries of the global North are right to assert that the development of thriving markets for copyrighted content is a desirable goal, but their resistance to access reforms designed to promote educational capability is shortsighted. As Margaret Chon puts it, "Education is fundamental to the capacity-building upon which all further progress is made."⁸ And as Ruth Okediji notes, in developing countries "the capacity to infringe is significantly limited by the lack of computers to access online content."⁹ The neoliberal view of the primacy of markets, which predominates within the international intellectual-property system, regards strong intellectual-property protection as an essential precondition for cultural and technological progress. For many developing economies, that theory puts the cart before the horse. Bridging the digital and educational divides that separate developing and developed countries requires sustained investments in literacy and capacity-building. Without the broadly distributed capability to participate in the development of intellectual goods, some countries simply will not develop intellectual property industries at all. Gross inequalities in resources and education levels cannot simply be assumed away; such inequalities also must be starting considerations in the design of a global copyright regime.

A final important area for access reform concerns provision of the infrastructure necessary to enable situated subjects to enjoy the fruits of copyright liberalization. In an increasingly digital world, enjoying access to cultural resources requires access to information networks. The A2K movement has been instrumental in connecting intellectual property issues to network access and architecture issues, reminding us that access to knowledge has dimensions that reach beyond intellectual property law. More specifically, the net-neutrality movement emphasizes the important role that nondiscriminatory interconnection policies play in securing access. In developed countries, the principal vehicle for addressing access and nondiscrimination issues is national broadband policy. The leading empirical study of national broadband policies, conducted by Harvard's Berkman Center for Internet and Society at the request of the U.S. Federal Communications Commission, concludes that a range of policies designed to promote open access at various stages of the communication process "are almost universally understood as having played a core role in the first gen-

eration transition to broadband in most of the high performing countries . . . [and] now play a core role in planning for the next generation transition.”¹⁰ In developing countries, where infrastructures are lean to nonexistent, and where the resources for large-scale broadband deployment typically do not exist, policy makers and entrepreneurs have experimented with an array of other mechanisms for providing network access, including satellite transmissions and mobile-communication platforms.

Access to cultural artifacts, however, does not necessarily go hand in hand with either the legal right or the technical ability to use them as creative practice requires. Some legal scholars argue that users enjoy more latitude than a reading of the copyright law would suggest. In particular, they assert that content owners’ increasing acceptance of “remix culture” creates a category of “tolerated use” upon which users can safely rely.¹¹ That may be correct, but it is far too early to proclaim a stable equilibrium. Controlled “playgrounds” such as the one described earlier in this chapter offer corporate intellectual-property owners a new direction for tolerated use, and one that does not inevitably afford broad user freedoms. In addition, new models for cloud storage and the authenticated delivery of media content may foreclose some types of use as a technical matter. Some scholars argue that because rights of use and reuse are necessary for cultural and political participation, access should be conceptualized as including such rights.¹² These scholars are right to think that not all kinds of access are equally valuable; for exactly that reason, however, an access principle alone cannot provide a sufficient foundation for a just information policy. The need to afford latitude for situated users signals the need for a different and complementary principle, which relates to what I have called semantic discontinuity, or interstitial flexibility within the system of legal rights, institutional arrangements, and associated technical controls.

In addition, rights of access to information and information networks do not necessarily correlate with rights to privacy; indeed, they more typically function in the opposite way. As network users become habituated to trading information for information and other services, access to goods and services takes place in an environment characterized by increasing amounts of both transparency and exposure. Similarly, access to the network and to networked information artifacts can be a double-edged sword, depending on how the network and its constituent artifacts are configured. To an increasing degree, features of the networked information environment are characterized by seamless opacity and highly granular authorization processes. These changes have implications for two clusters of important social values, one relating to transparency and accountability and the other relating to informational and spatial freedom of movement. Neither set of concerns can be addressed effectively simply by mandating access to information goods and services. Instead, human flourishing in the networked information society requires additional structural safeguards.

Operational Transparency

To benefit fully from the opportunities that networked information technologies offer, network users require not only access to information delivered by networks and devices but also access to information about the way those networks and devices function. I will use the term “operational transparency” to refer to three distinct sets of transparency concerns, each of which relates in a different way to the technical and informational operation of networked processes and to the effects of those processes as experienced by situated users. Operational transparency encompasses transparency about the design and implementation of surveillance practices, transparency about the operation of the network’s borders and flows, and transparency about the processes by which network standards are designed and adopted.

The rationale for identifying operational transparency as a condition of human flourishing in the networked information society is straightforward. The lives of situated subjects are increasingly shaped by decisions made and implemented using networked information technologies. Those decisions present some possibilities and foreclose others. Most people have very little understanding of the ways that such decisions are made or of the options that are not presented. In many cases, this facial inaccessibility is reinforced by regimes of secrecy that limit even technically trained outsiders to “black box” testing. We would not tolerate comparable restrictions on access to the basic laws of physics, chemistry, or biology, which govern the operation of the physical environment. The algorithms and protocols that sort and categorize situated subjects, shape information flows, and authorize or deny access to network resources are the basic operational laws of the emerging networked information society; to exercise meaningful control over their surroundings, people need access to a baseline level of information about what those algorithms and protocols do.

The first category of operational transparency issues relates to the surveillance processes employed by both public and private entities to sort and categorize individuals and groups. Existing regulatory frameworks that have attempted to ensure transparency about the collection and use of personal information do not fully address the need for operational transparency in surveillance. In the United States, most private-sector uses and transfers of personal data are regulated only by the Federal Trade Commission’s general prohibition against unfair and deceptive trade practices. Most reputable firms that deal directly with consumers do disclose some information about their “privacy practices,” but the incentive is to formulate disclosures about both purposes and potential recipients in the most general terms possible. This practice in turn shields secondary recipients of personal data, many of whom do not disclose information about their activities at all. In contrast, the guidelines on fair information practices adopted by the Organization for Economic Cooperation and Development (OECD) and enacted as a directive by the European Union require parties that collect personal information to provide disclosures that specify the purposes for which the information will be used and any potential recipients other than the original collector. They also must afford a meaningful opportunity to examine and correct the information.¹³

From the standpoint of operational transparency, even the more stringent regulatory regime adopted in the European Union has two major defects.

First, the statement of fair information practices, which dates from the 1970s, seems to envision a series of discrete information-collection events, and is poorly suited to the reality of dynamic information-collection in online environments. More fundamentally, even highly granular, dynamically-updated purpose and recipient disclosures would not necessarily shed light on the operational significance of collected information. Telling someone what pieces of information were considered for the purposes of making decisions about credit or medical coverage or targeted advertising provides no information about how that information mattered, about the other assumptions used to construct the operational heuristic, or about how different information would have changed the result. Yet such operational information makes all the difference, and liberal legal thinkers should have no trouble understanding why. Operational disclosures are essential both for informed consumption of the goods and services that are offered and for open, informed debate about the processes by which individuals and groups are sorted and categorized. In some cases, such disclosures might lead the affected individuals to change their own behavior; in others, they might lend concrete support to calls for regulatory reform.

Internationally, pressure to strengthen fair-information-practice guarantees is mounting. In 2011, the OECD began considering whether the fair-information-practice guidelines require revision to reflect the realities of modern data-processing practice.¹⁴ At minimum, such reforms should address the problems of dynamism and operational significance. For a guarantee of transparency to be meaningful, people who are the subjects of information processing need enough information to enable them to understand more accurately both how items of information in their own profiles will be used and more generally how particular types of decisions are made.

A second category of operational-transparency concerns relates to information about the network's geographies of accessibility. As we saw in Chapter 8, those geographies are increasingly opaque to network users, concealed by the seamless operation of autonomic technologies. Most ordinary network users have little understanding of how networked information processes work, yet those processes mediate access to an increasingly broad range of public and private services. Danielle Citron's important work on "technological due process" explores the use of networked information technologies to perform a variety of public functions traditionally associated with government, ranging from the tabulation of election results to eligibility assessments for benefits. Citron argues persuasively that guarantees of due process that apply to government action should be extended into the realm of network architecture.¹⁵ A regime of technological due process would require the public provision of meaningful information about the ways that traditionally public functions are performed, and would extend that requirement to the capabilities of technical systems supplied by private contractors. In addition, it would increase accountability by imposing strict audit requirements and authorizing legal challenges to inadequate disclosures.

As we have seen, though, operational-transparency requirements that apply only to the public sector are not enough to achieve operational transparency of the network as experienced by situated users. Geographies of accessibility and inaccessibility are comprehensively mediated by proprietary platforms and algorithms. Frank Pasquale has proposed that search engines be sub-

jected to a regime of “qualified transparency,” which would mandate disclosure of their practices for filtering and displaying search results. As Pasquale explains, such a regime need not involve unrestricted public disclosure of trade-secret information in order to effectively convey to individuals the information that they need.¹⁶ Various devices are available to protect commercially important secrets from general disclosure, including disclosure to an ombudsman or expert panel that would verify the accuracy of publicly reported operational information. By analogy, policy makers could design similar regimes of qualified transparency for the other technologies discussed in Chapter 8, such as nascent rights-management systems that pair cloud storage with plug-and-play portability, or the regimes of unremarkable computing envisioned by designers of networked artifacts.

The final category of operational-transparency concerns relates to the processes by which general network standards are developed and implemented. As we saw in Chapter 8, both secrecy and technical mystification tend to frustrate sustained public scrutiny of standards processes and their outputs. Some of those standards, such as the basic Internet Protocol, are matters of great public concern. For such standards, the public interest in operational transparency is not satisfied by providing for qualified transparency after standards have been determined. In her study of the political struggles surrounding the development of IPv6, Laura DeNardis concludes that best practices in Internet-standards governance require commitment to a standard that is “open in its development, open in its implementation, and open in its use.”¹⁷ A policy of nominal openness toward interested and technically skilled participants is not enough to satisfy the first criterion. Even bodies like the Internet Engineering Task Force, which has maintained a strong commitment to open standards development, must confront “intrinsic barriers to participation related to technical expertise, language, funding, and culture.”¹⁸ Adequate operational transparency requires the design of participatory mechanisms that take those barriers into account. To achieve openness in implementation, DeNardis recommends open publication of and royalty-free access to general network standards. A well-designed regime of qualified transparency, tasked with identifying and disclosing departures from such standards, would work to ensure openness in use.

This preliminary analysis of operational-transparency interests leaves unanswered important questions about the dividing line between those technical protocols that should be subject to open-process requirements and those that should be subject only to the less stringent requirement of qualified transparency. Strategies for designing both types of regimes, and for generating increased commitment to open development processes, are important subjects for future research and experimentation. Ideally, a well-designed regime of qualified transparency would exert upward pressure on development practices, bringing to the public’s attention the existence and capabilities of proprietary protocols that shape networked information processes and generating systemic pressures toward even greater disclosure.

Even high levels of operational transparency, however, would not necessarily equip networked individuals and communities with the resources to counter continuing interstitial shrinkage in the legal and institutional frameworks that regulate access to and use of information and information networks. Imagine that existing delivery systems for copyrighted content have been re-

placed by a system for cloud storage and automated delivery to users wherever they are located via a variety of fixed and mobile devices. The pricing structure is clear and uncomplicated, and the delivery mechanisms are effective and largely bug free. The delivery system incorporates clearly disclosed technical rules that authorize the consumptive use of media content for which access fees have been paid, but that do not authorize copying for other uses. Or imagine a clear rule stating that a provider of financial services may purchase whatever information it wants about individuals, may incorporate that information into its algorithms governing credit availability and pricing as it sees fit, and may sell the information to whoever it wants. Individuals are told exactly how their transactional and personal histories will affect their eligibility for various financial services and interest rates.

In both of these examples, the rules are clear and transparent—and in that respect, they represent some improvement over current practice in both areas—but the interstitial flexibility that they afford for the processes of creative practice and subject formation is minimal. These examples illustrate that, like access, operational transparency may be a necessary condition for human flourishing in the networked information society, but it is not a sufficient condition. The examples of what operational transparency does not cover point to a third structural condition for human flourishing, which relates to whether the legal and technical rules that govern access to and use of information resources contain sufficient interstitial complexity to facilitate cultural mobility, the boundary-making processes of privacy, and the play of material practice.

Semantic Discontinuity

In an age characterized by increasingly seamless and granular regulation of information access and use, and by increasingly precise efforts to monitor and predict individual behavior with comparable seamlessness and granularity, preserving adequate room for play within the domains of culture, subject formation, and material practice requires regulatory and technical interventions designed to foster what I will call “semantic discontinuity.” Semantic discontinuity is the opposite of seamlessness: it is a function of interstitial complexity within the institutional and technical frameworks that define information rights and obligations and establish protocols for information collection, storage, processing, and exchange. Interstitial complexity permeates the fabric of our everyday, analog existence, where it typically goes unappreciated. Its function, however, is a vital one. It creates space for the semantic indeterminacy that is a vital and indispensable enabler of the play of everyday practice.

Systems of logical reasoning that derive from the tradition of Enlightenment rationalism have enormous difficulty in acknowledging the importance of semantic discontinuity. Instead, instances of semantic discontinuity tend to be conceptualized as imperfections that detract from the realization of legal, market, and technical ideals. More specifically, in thinking about the optimal regulatory structures for the networked information society, law- and policy makers are caught in a tug-of-war between two logical principles that together operate to deprive the networked self of the shelter afforded by interstitial complexity. The first principle holds that a system of law should neither draw nor respect lines between different types of conduct unless those lines can be justi-

fied by reference to first principles. A canonical statement of this principle is Oliver Wendell Holmes's injunction that the role of the law is to make rules for the "bad man," who will exploit semantic or technical loopholes in any way that he can.¹⁹ The second principle holds, at times paradoxically, that a system of law governed by principles of rationalism and provided with enough information can derive the right system of rules to govern individual conduct.

To a considerable extent, it is the push and pull between the logician's skepticism and the technocrat's confidence that produces the dynamic that this book has explored. Unwilling to acknowledge a potential conflict between copyright and creativity, we demolish assertedly arbitrary boundaries that limit copyright scope while at the same time insisting on the possibility of defining properly tailored limitations and exceptions to shelter the lawful activities of users. Unwilling to acknowledge a potential conflict between unlimited flows of personal information and sound social and economic policy, we eliminate seemingly arbitrary restrictions on the collection, processing, and exchange of personal data while at the same time insisting on the possibility of protecting autonomy by means of narrowly defined zones of personal privacy in appropriate cases. In so doing, we implicitly accede to technocratic regimes of private governance organized around the processing of personal information. Together, these attitudes produce a schizophrenic approach to the design of network architectures. In some cases, we prize technical openness; in others, technical closure.

As we have seen, the continuities of information flow imposed by emerging market and legal institutions signal institutional realignments that are not simply logical or technical, but also and more fundamentally political. Those realignments enhance the power of actors and institutions that benefit from commercially continuous flows of cultural and personal information. At the same time, they diminish the ability of individual users and communities of practice to encounter and interact with flows of culture, and to pursue contextually specific practices of self-definition, in patterns that form and re-form more organically. Powerful actors that benefit from emerging regimes of authorization argue that it is unrealistic to expect the rhythms of digitally mediated life to match earlier ones; users who want the freedoms that new communications technologies bring should be prepared to make some sacrifices. But that argument confuses two kinds of inevitability: the fact that emerging patterns of information flow serve powerful economic and political interests, and thus might have been predicted by anyone paying attention to the distribution of incentives, does not make the patterns natural or just.

A commitment to human flourishing in the networked information society requires an effort to reverse, or at least cabin, the tendencies toward seamless continuity within legal, market, and technical infrastructures for information exchange. In making law for the bad man, we must not lose sight of the fact that we are also making law for the good person and for the good society more broadly. Creativity, critical subjectivity, and everyday practice flourish in conditions of (partial) unpredictability, and humans require creativity, critical subjectivity, and everyday practice to flourish. With practices this foundational at stake, freedom can consist in privileging discontinuity for its own sake.

Participation within Cultural Landscapes

Chapter 4 argued that copyright law- and policy making should foreground the everyday practice of situated users and the constrained yet open-ended process of working through culture. Copyright creates the legal foundation for capital investment in cultural production, and that function is an important one both economically and culturally. At the same time, culture requires room to move, and its movement benefits all who participate. Copyright should be understood as striking a balance not between present authors and the abstract “public” but rather between the near-term goal of creating economic fixity and the longer-term goal of fostering cultural mobility. So understood, the system of copyright requires the deliberate introduction and maintenance of legal and institutional discontinuities that shelter cultural play.

Copyright law purports to recognize a discontinuity principle, but increasingly pays only lip service to it. The contemporary approach to defining copyright rights assigns owners broad and often overlapping rights and then subjects those rights to narrow, situation-specific exceptions and limitations. The formulation of rights reflects what I have described as the logician’s skepticism; it is deeply suspicious of approaches that might exclude new modes of expression or foreclose newly developing markets. In the realm of exceptions and limitations, however, law- and policy makers become mysteriously confident in their ability to define criteria for unremunerated uses in precise, granular terms, without regard to the inherent unpredictability of future events that might threaten such uses. Taking their guidance from these attitudes, courts interpret rights in an expansive fashion and construe exceptions narrowly. Both halves of this approach reflect a near allergy to logical gaps in copyright’s coverage. Copyright rights have a protean quality, expanding into every avenue of potential profit. With one significant exception, copyright limitations generally have not demonstrated a parallel capacity to evolve as technologies change.

Copyright scholars and A2K advocates have advanced a variety of proposals for recalibrating the copyright balance. Those proposals generally fall into four categories. First, many argue that courts can and should interpret the fair use doctrine more expansively, developing a jurisprudence that would more predictably privilege a broader variety of uses. Second, advocates of fair use reform often argue that identifiable communities of practice—documentary filmmakers or fanvidders, for example—should develop statements of “best practices” that judges might consult for guidance. Third, copyright scholars have proposed statutory reforms to copyright formalities intended to speed the transfer of certain types of copyrighted content to the public domain. Fourth, the Creative Commons movement and similar open-access movements have sought to encourage more widespread use of open licensing regimes by which copyright holders can permit various uses of their works that satisfy particular criteria. All these proposals represent very good ideas. Without more, however, they are unlikely to yield the sort of meaningful recalibration that creative practice requires.

Within the U.S. legal system, the fair use doctrine has been the important exception to the general consensus that limitations on copyright must be narrow and precisely defined. Fair use shelters some uses in a way that purports to be open ended and sensitive to the ultimate value of the use in question. In practice, however, the shelter that fair use affords for everyday practice is not

nearly as capacious as users tend to assume. Businesses and courts uncomfortable with open-ended exceptions have developed interpretive rules designed to make fair use more manageable. Those rules have operated to constrain user privilege within relatively narrow channels. They privilege uses that are clearly identifiable as criticism, parody, or biography, but do not reliably privilege many other forms of reworking that are central to contemporary creative practice, including allusion, homage, and pastiche, nor do they reliably privilege copying that situated users might undertake for other, less directly expressive reasons. In cases involving the literal copying of small excerpts—for educational use, for inclusion in a documentary, or for any number of other purposes—a practice has emerged of licensing “clearances.” Courts cite these emerging licensing markets as justifications for rejecting fair use arguments, producing a cycle of “doctrinal feedback” in which the zone of fair use continually shrinks.²⁰

Although some aspects of fair use practice are probably amenable to reform, this history suggests that fair use likely cannot function as the general-purpose exception that some A2K advocates have envisioned. Even if courts took steps to end the problem of doctrinal feedback, the culture of licensing likely would persist because that culture is not only or even primarily a judicial creation. The culture of licensing is first and foremost a risk management culture; it is the by-product of culture facilitators’ demand for clear rules with which to structure their dealings. Best-practice statements that are clear, concise, and aggressively promoted can counter that demand to some extent, but they confront a chicken-and-egg problem; to work most effectively, they must interrupt licensing norms at their inception. In addition, the best-practice model requires a community that is sufficiently well established to begin with. For the ordinary situated user of copyrighted content, who has only what Larry Lessig describes as “the right to hire a lawyer,” best-practice statements offer little realistic possibility of shelter.²¹

Proposals for the “reformatization” of copyright and the adoption of open licensing norms, meanwhile, would simply reinforce the relative power of mass commercial culture. Consider, for example, the proposal that United States revert to a rule requiring renewal of copyright after an initial fixed term, or the proposal that it reinvigorate copyright registration and notice rules by making compliance with formalities a condition of claiming enhanced remedies.²² To an overwhelming degree, copyright owners of mass commercial culture would take advantage of such rules, while many individual proprietors of copyright would not. Similarly, open-access initiatives such as Creative Commons have achieved very little penetration within the core domains of mass copyrighted culture, and those regimes often impose other transaction costs of their own.²³ Under either proposal, therefore, the common cultural baseline established by mass commercial culture most likely would remain off-limits to many forms of creative play. That result does not serve situated users’ need to interact meaningfully with the constituent elements of the cultural landscape that surrounds them.

A different kind of strategy for privileging certain kinds of use would involve defining the copyright rights themselves in a more limited fashion. When confronted with this possibility, most copyright lawyers and scholars automatically resist it. If pressed, some will concede that copyright owners

would not be materially worse off than they are now if they enjoyed only the right to control commercial exploitation of their works. Unremunerated but commercially harmful uses, such as P2P file sharing, could be addressed by defining the rights to include any exploitation that has significant commercial impact. Similarly, some U.S. copyright lawyers are willing to entertain the possibility of reformulating the derivative-work right as an adaptation right as long as it includes language that clearly signals the reservation of reasonably expected and commercially significant adaptations to the copyright owner. Commercial significance, though, is one of the most notoriously elastic concepts in copyright jurisprudence; it rapidly expands to cover everything in sight. Even in cases where evidence of commercial harm is deemed speculative, courts usually think it is important to leave the door open to a finding of harm later.²⁴ We might therefore be reasonably confident in predicting that rights defined by reference to commercial expectations, and unconstrained by any definite outer boundary, would be subject to continual judicial expansion.

In short, copyright lawyers and policy makers can talk the talk of limiting rights in the interest of “balance,” but when it comes to actually ratcheting back the scope of copyright rights in any meaningful way, they become oddly reluctant to make changes that might affect the copyright owner’s future bottom line—even though they can’t predict whether there will actually be an effect or how much effect there might be. One reason for this reluctance, which is well understood by copyright scholars, is what we might call a naive restitutionary impulse—the idea that commercial gain to anyone else constitutes an injury that demands compensation so the copyright holder can be made whole.²⁵ If we need to adopt a broad reading of one or the other of the copyright exclusive rights to do this, then we should do it. However, the reluctance to limit rights does not depend entirely, or even principally, on the restitutionary impulse.

The more fundamental reason that copyright judges and policy makers resist setting meaningful limits has to do with the form of reasoning that our legal culture prizes most highly. It is best illustrated with an anecdote. Upon occasion, I have asked groups of upper-level law students to describe their exam-taking strategy. I ask them to imagine that they are taking an exam in some other, non-copyright-related subject—torts, or maybe constitutional law—and that they have been presented with a long, complicated fact pattern and been asked whether the plaintiff can succeed on any of a number of theories of relief. Then I ask whether they think they would get better grades by arguing that the plaintiff should succeed or fail. Except for the odd contrarian, who seeks to stand out by going against the trend, the students tend to believe that unless the professor has signaled a clear preference for a different strategy, they will get better grades by attempting to show how the plaintiff could succeed, even if it required an expansion of the grounds recognized by the law as basis for recovery.

That my students think the way they do about exam performance is no accident. We—I and my colleagues in law teaching—have taught them to think that way. To some extent, their response reflects successful internalization of the common-law method, in which the definition of legal rights proceeds by flexible incrementalism. But my students also understand arguments for extension as demonstrating both more technical creativity—more skill at lawyering—and more “true” understanding of the subject matter. You “really” under-

stand torts or constitutional law or whatever when you can explain why a particular right “really” extends to cover situations in which it has never before been applied. Within the U.S. legal system, at least, the extension of rights and remedies into new territory is the essence of what lawyers do, and skill at it is a key indicator of professional and intellectual excellence.

That belief system, in turn, is both an intellectual legacy of Enlightenment rationalism and a more direct descendent of twentieth-century legal realism. The understanding of rights as narrow entities with fixed limits has become identified with the irrationality of nineteenth-century classical legal thought and the formalist categories on which it relied. At its inception, the realist project was both philosophical and political; it sought to show that limited understandings of legal rights and obligations served fundamentally antiprogressive ends. Over time, however, the reformist impetus waned and was replaced by a positivist antiformalism.²⁶ Within the context of the original realist project, it made particular sense to interpret limits on rights and remedies as arbitrary fictions propped up by a discredited set of intellectual and political commitments. Positivist antiformalism affects a more neutral stance, resisting doctrinal distinctions that appear to have no foundation in logic. The 1976 Copyright Act and most other contemporary copyright laws and treaties are realist documents in this latter sense: they subordinate careful consideration of the balancing problems involved in copyright policy to the goal of drafting rights in a way that avoids artificial constraints. Within U.S. copyright circles, to suggest deviating from a baseline of broad entitlements subject to narrow limitations is to evoke the much-reviled categorical structure of the law that the 1976 Act replaced, the 1909 Copyright Act, which doled out narrower rights in an ad hoc, historically contingent fashion.

The limits of a legal methodology that treats limited rights as logically and intellectually disrespectable are particularly evident in contexts in which it is necessary to balance competing, equally important interests. If we think about the patchy, incomplete structure of the 1909 act from the perspective of the creative process, its logically discontinuous structure was a feature, not a bug. The 1909 act was a product of its time, a deeply formalist text whose authors understood, for example, “lectures” and “essays” as performing lexically different functions. That structure reflected a way of thinking about creative works that cannot, and likely should not, be recaptured. But principled insistence on real limitations is an intellectual stance that deserves better treatment. Creative practice flourishes most fully under conditions that permit unexpected encounters with new information and that provide room for reworking, tinkering, and other forms of creative play. As we saw in Chapter 6, the forms of creativity that we prize in artistic and intellectual endeavor, and that we cite when talking about what we expect the copyright system to protect, often turn out to depend centrally on imitation and reworking. From the perspective of creative practice, a legal regime characterized by formally incomplete rights—by logical gaps that permit at least some uncontrolled access and use—comes closer to solving the balancing problem that copyright confronts.

The 1909 act thus suggests a more general template for achieving semantic discontinuity in copyright law and practice: a copyright regime concerned with the balance between economic fixity and cultural mobility should replace broad, all-encompassing statutory provisions and generous, judicially

created tests for infringement with narrower, clearly delimited formulations of copyright rights separated by deliberate gaps representing uses that are reserved in the service of cultural play, regardless of commercial consequence.

First, a regime of copyright recalibrated to prize and facilitate play should create broad zones of what Jessica Litman has called “lawful personal use”—uses of copyrighted works that are reserved to the situated user regardless of whether regimes of authorization might be developed to monetize them. Zones of lawful personal use should be defined using the sort of broad statutory language that until now has been reserved for the definition of copyright rights. A statutory provision for lawful personal use also should avoid subjecting such uses to a rigid form of the public-private distinction, such as the one that qualifies the current statutory definition of public performance. Many personal uses do occur in private spaces, but personal uses do not occur only at home or among a family and its circle of social acquaintances. They occur at work, at school, on trains and airplanes, and in many other places. Here copyright can usefully draw lessons from recent privacy scholarship’s emphasis on the importance of context: like privacy, lawful personal use is subject to norms of contextual appropriateness and flow; copyright should recognize and respect those norms.²⁷

Second, a regime of copyright recalibrated to facilitate play should sharply limit copyright owners’ rights to control adaptations and remixes by third parties. It should do so by enumerating, in a fashion designed to emulate the 1909 act’s discontinuous style, the list of adaptations to which copyright-holder control is permitted to extend. The list should not simply ignore fanworks and other examples of remix culture, as current copyright law now does, but should clearly reserve a broad range of remix privileges to users. Elsewhere I have written about how copyright law might apportion sequelization rights, reserving for authors the right to develop for commercial exploitation continuations of the original story while permitting users to experiment with fanworks and also to develop for commercial exploitation new stories set in other authors’ fictional worlds.²⁸ In other cases—derivations of visual artworks or musical works, translations, and so on—the law should specify a similar apportionment and should instruct courts to safeguard the interests of users when deciding cases that fall near the boundary.

Third, in a regime of copyright recalibrated to facilitate play, rules about the privileges and obligations of information intermediaries should be designed with the balance between fixity and play in mind. They should facilitate, rather than impede, the efforts of intermediaries such as digital libraries and search engines to organize information and present search results in ways that are useful. The Google Book Search project is a paradigmatic example of the sort of use that should be permitted to both Google and would-be competitors because of the extraordinary social benefits it will produce. In addition, copyright law should set strict limits on indirect liability and technical-protection rules that cast a chill over the development of new technologies, including user-driven innovations. As we have seen, the problem is not simply that technology developers risk infringement liability for giving users too much flexibility, but also that the ongoing, legally mediated realignment of technology markets encourages technology providers to overcompensate in the opposite direction, giving users too little flexibility. A properly constructed regime

of indirect liability would seek to minimize both risks by explicitly inserting the legitimate interests of users into the liability calculus.

Reforms such as these, designed to restore interstitial complexity within the copyright system, would equip situated users of copyrighted content with the legal standing to direct flows of culture according to their own legitimate expectations and needs, and thereby to participate in the culture evolving around them. For exactly that reason, such reforms also would go a long way toward curing the legitimacy crisis that the copyright regime currently confronts. Although that crisis has resulted in part from market and enforcement behaviors that users perceive as abusive, it is also fueled by the copyright industries' habitual practice of normative overclaiming. Situated users are told that copyright is itself the font of all creativity, even as they daily experience that claim to be false. Clearer acknowledgment of limitations on what the copyright system can claim for itself would produce greater public respect for those claims.

Boundary Management within Social Landscapes

Chapter 6 developed a working definition of privacy as the process of differential boundary management by situated subjects and argued that privacy so defined is an indispensable enabler of the process of subject formation. In the case of privacy, the fixity that threatens emergent subjectivity is bound up with constitutive ideologies about the relationship between information processing and truth, and about the primacy of individualized treatment. While the ability to identify individuals persistently and accurately is important for some purposes, a just regime of information policy also must seek to provide the breathing room that critical subjectivity requires. Contextual integrity requires interstitial complexity; privacy law and policy should reinforce and widen gaps within the semantic web so that situated subjects can thrive.

Although privacy law purports to recognize a discontinuity principle, that principle operates primarily to protect small islands of concededly "intimate" or "sensitive" information and correspondingly small enclaves of acknowledged physical seclusion. In an age of distributed information processing, moreover, even those islands are fast eroding. As we saw in Chapter 3, discourse about surveillance practices and privacy laws operates primarily in the mode of technocratic confidence. Both private and public actors believe that if we can just collect enough information about people, the route to enlightened decision making in the realms of both profit and policy will be revealed. The information-processing imperative dictates that if more information can be collected and incorporated into predictive profiles and algorithms, it should be. We can be deeply troubled by particular uses of information but still believe, equally deeply, that the cure for misuses is even more and better information. Disrupting the information-processing imperative violates an implicit equation that is fundamental to the paradigm of liberal political economy: information equals truth.

Discourse about privacy protections, meanwhile, operates in the minor key of logical skepticism. The information-processing imperative has an important corollary, which I will call the "Luddism proviso": predetermined limits on information processing are a manifestation of irrationality, and those who endorse them are fundamentally antiprogress. Some kinds of information are

more private than others, but to the extent that privacy protection for particular items of information can be made to appear arbitrary—because the information is relevant to a contracting party’s decision to provide services, because its sharing might enable valuable efficiencies, or because it has already been disclosed to somebody else anyway—existing legal restrictions begin to fall away, and new, more effective privacy protections fail to materialize.

Consider, for example, the privacy provisions of the Gramm-Leach-Bliley Act, which governs the collection and exchange of information by U.S. financial institutions. Some legislators and privacy advocates favored restricting the extent to which such institutions could share customers’ personal information with both affiliated and nonaffiliated companies. Lobbyists argued that such a rule would raise costs to firms seeking to market, and consumers seeking to comparison shop for, financial services. The bill reported from committee in the House of Representatives included an opt-out rule covering both affiliates and nonaffiliates. As finally enacted, the opt-out rule covers only information sharing with nonaffiliates. It permits information sharing with affiliates without limitation, on the implicit presumption that information given to one member of a corporate “family” isn’t private as far as the other members are concerned.²⁹ Since both vertical and horizontal integration are widespread in the U.S. financial services industry, this rule facilitates an enormous amount of information sharing.

Next, consider the “deidentification” standard promulgated under the Health Insurance Portability and Accountability Act (HIPAA), which authorizes disclosure to third parties of data that includes birth year and partial zip code information. When combined with information from other readily available data sources, such entries can be reidentified with relative ease.³⁰ Interested parties have resisted additional restrictions on data disclosure, citing the need to conduct accurate population research. The specific nature of that need is left strategically vague, and the larger structure of privacy discourse allows it to remain that way. From a social welfare standpoint, not all needs are equal. We might think, for example, that medical researchers have greater need for population data than marketers of personal-care products do. To that extent, the drafters of the HIPAA rules agreed with the unequal-value proposition; medical researchers may acquire fully identified data if they observe other confidentiality requirements.³¹ Yet they do not appear to have considered seriously whether some “population research”—for example, research designed to identify the population of potential customers for adult incontinence products, or research designed to develop more precise differential health-insurance pricing—has so little social value that we should not worry unduly about frustrating it.

Privacy scholars and advocates have advanced a variety of proposals for more effective protection, but proposals to impose substantive limits on information processing and sharing tend to make A2K advocates uncomfortable. Many U.S. legal scholars and technology commentators have tended to think that privacy problems can be addressed in a less heavy-handed fashion by giving people access to “privacy-enhancing technologies,” such as services that enable anonymous Internet browsing, and by providing more comprehensive information about the privacy practices of public and privacy entities.³² As the Gramm-Leach-Bliley and HIPAA examples illustrate, however, neither palliative affords meaningful shelter in the thousands of everyday contexts in which

one must supply personal information in order to engage in ordinary transactions and receive important services.

But the liberal-rationalist tradition is not solely to blame for the dominance of the view that equates individualized information with truth and anti-individualization rules with antiprogressive animus. Liberal political theory's discourse of rights and human dignity also emphasizes individualized treatment at moments of decision. Privacy scholars sometimes draw explicit contrasts between aggregated treatment and individual dignity; thus described, the problem is not simply that invasions of privacy objectify individuals, but that they do so in a way that denies individuality itself, filtering out potentially individuating elements of context and subjecting individuals to categorical judgments. Yet that mode of reasoning about privacy contains the seeds of its own undoing. When the dignity interest is formulated in terms of a right to individualized treatment, it becomes difficult to argue that making health coverage or financial decisions based on highly granular profiles is fundamentally unjust.

The reluctance to make normative decisions about the limits of information processing is not well founded. First, a wealth of historical evidence undercuts the rationalist faith in the inevitable link between information processing and truth. Innovations in information processing are not invariably linked to just and wise social policy. Automated census technologies have been used to facilitate persecution and genocide, and automated surveillance technologies to support regimes of political repression.³³ Law- and policy makers tend to understand these examples as instances of conceptually unrelated (and morally repugnant) social ideology run amok. On that understanding, the problem is that bias occasionally diverts rationalism from its true course. But distinguishing rationalism from bias requires an omniscience that situated policy makers do not possess.

The insistence that dignitary concerns inevitably require individualization is equally curious. In the era of automated personalization, we have come to realize that individualization is not a sufficient condition of dignified treatment. The new personalized information services enabled by the semantic web are highly individualized, but still make judgments in formulaic and sometimes objectifying ways. Yet we have continued to act as though individualization is a necessary condition of dignified treatment. Scholars who study the moral dimensions of profiling argue that the notion of individualized treatment is inherently slippery because we cannot avoid inferring individual characteristics from group attributes.³⁴ On that account, the notion of perfectly individualized treatment is fictive, an unattainable ideal. It seems more sensible to inquire whether, both practically and theoretically, the ideal of individualized treatment simply cannot support the normative weight it has been asked to bear. Critics of the individualist tradition in liberal jurisprudence have long argued that it devalues other strands in our moral tradition that are predicated on equality and that are emphasized within the capabilities approach. In a world of increasingly ubiquitous information processing, perhaps the theory of privacy as room for boundary management can help point the way toward a different way of thinking about the requisites of human dignity.

If individualized treatment is not necessarily dignifying, perhaps the reverse is also true: perhaps dignifying treatment is not necessarily individualized. Put differently, if individualized treatment can be dignifying or objectify-

ing, perhaps the same is true of aggregated treatment. Decisions affecting individuals and groups within society can be classified as individualized or aggregated, and as dignifying or objectifying, but the pairs of attributes need not always align in the same way. Legal and policy decisions affecting individuals and groups can be conceptualized using a matrix (below) that allows for the possibility of actions that are both aggregated and dignifying. Within privacy law and theory, as within liberal political theory more generally, the lower right quadrant of the matrix is undertheorized. It is occupied principally by the claims of various identity groups to equal protection of the law, but its potential extends far beyond such claims. An important purpose of privacy law and policy is to populate that quadrant of the policy matrix, advancing the concept of dignifying aggregation in a way furthers a more general, non-identity-based claim to the right to develop capabilities for human flourishing.

	Objectifying	Dignifying
Individualized	Profiling; Semantic Web	Due Process
Aggregated	Bureaucracy; “One-size-fits-all”	Equal Protection; Privacy

So conceptualized, a “just aggregation” principle underwrites an equality-based right to avoid individualized treatment, including both practices aimed at transparency and practices aimed at exposure. The situated subject requires protection against information-processing practices that impose a grid of highly articulated rationality on human activity, and against the reordering of spaces to institute norms of exposure and collective objectification. Against a background of increasing convergence, effective legal protection for privacy requires interventions aimed at preserving the commercial, technical, and spatial disconnects that separate contexts from one another. Policy interventions designed to promote semantic discontinuity should operate both informationally, by disrupting the grid, and spatially, by affording shelter. And on this understanding of privacy’s purpose, privacy consists in the setting of limits precisely where logic would object to drawing lines.

A regime of discontinuity-based privacy protection informed by a just-aggregation principle would set stringent limits on the collection, use, retention, and transfer of personal information. Such restrictions in the United States too often reflect a purely proceduralist conception of consumer protection that revolves around notice and consent. Many other countries provide more meaningful protection, but have struggled to enforce privacy guarantees against data processors located outside their borders. For privacy protection to be effective in preserving room for emerging subjectivity, privacy guarantees must be substantive and global, aimed at introducing more than modest amounts of interstitial complexity into the semantic web.

For most cases involving use of personal information by commercial and nonprofit entities, a data fiduciary model based on fair information practices establishes a baseline standard for protection.³⁵ Some aspects of this protection might be waivable, but the conditions of waiver would need to be strictly defined and highly granular so that waiver could not become a tool for routine evasion of privacy obligations. In the Gramm-Leach-Bliley example

discussed above, information sharing with both affiliates and nonaffiliates could proceed on an opt-in, unbundled (per-recipient) basis, subject to the operational-transparency requirements described earlier in this chapter. Where transfer is allowed, however, transferred information must be subject to strict purpose limitations, nonwaivable prohibitions on further transfer, and mandatory data-destruction rules. As Paul Ohm explains, some potential recipients of personal data are more trustworthy than others; to ensure compliance with privacy restrictions, privacy regulators should develop a system for certifying trustworthiness and prohibiting transfers to uncertified recipients.³⁶ The primary data fiduciary, meanwhile, would be subject to similar limits in its own uses of information. More particularly, either the law or implementing regulations would set restrictions on the types of personal information that the institution could use in making decisions about pricing and other terms of service.

Because privacy expectations and needs vary contextually, this basic structure would require modification for at least the following four special cases. First are contexts in which situated subjects' own well-being requires the collection, long-term retention, and more widespread sharing of individualized personal information. The most compelling example of such information is health-related information, which must be collected, kept, and often shared in order to enable successful treatment. In such circumstances, the governing law should waive data-destruction requirements and should permit data transfer as necessary for the effective provision of treatment, but should impose robust security requirements for information access and storage and should subject data custodians to periodic, publicly disclosed audits. Government entities that need to maintain and share certain permanent records, such as property records, benefits records, and judicial dockets, should be excused from data-destruction requirements, but should be required to redact designated categories of information before making the records available to the public.

The second special case involves databases used in research, including research on medical, public health, and social welfare issues. As noted above, the initial experience with data deidentification requirements has shown that data sets are much easier to reidentify than had been thought. In part, this results from the widespread public availability of partial data sets that can be cross-linked and correlated; many of the other proposals advanced here would reduce that availability simply because they would erect higher barriers to data exchange. In part, the ease of deidentification results from reliance on systematic and therefore predictable practices in the assignment of anonymous identifiers; in such cases, randomization of the assignment process would make reidentification more difficult.³⁷ The strength of the information-processing imperative, however, suggests that calls for the public release of deidentified data sets would remain strong and that incentives to develop new methods of reidentification would remain high. In some such cases, partial privacy protection can be achieved by introducing "noise" into data sets at a level that does not impair their utility. As long as that effect is reversible, however—a state of affairs that regulatory requirements requiring logical proof of concept tend to encourage—the likelihood of eventual reidentification is strong.³⁸ Ultimately, then, privacy policy makers must directly confront the extent of the stated need for accuracy. In contexts where accuracy is important—public health modeling for pandemic detection, for example—accuracy can be offset with confidentiality and security requirements. In other contexts, however, where there is no such compelling need, the balance should be struck

ling need, the balance should be struck differently. In particular, many for-profit acquisitions of population data for marketing and product research should not be permitted at all.

The third special case concerns disclosures of personal information via social-networking platforms. Regulation of the content of such disclosures would be difficult and ultimately counterproductive. Users derive important benefits from sharing the details of their lives with friends, family, and others. Yet users also derive important benefits from being able to establish and manage boundaries, and social-networking platforms have not met that demand with capabilities that remotely approach the context sensitivity that task requires. Part of the solution to that problem involves the tools used to designate disclosures for particular recipients, and will be discussed below. In addition, the privacy impact of social-networking disclosures can be minimized by implementing strict, nonwaivable rules governing commercial partnerships and the cross-linking of affiliated data services. Such rules should generally mirror the basic data-fiduciary model. Social-networking platforms should be able to share some information with trusted advertisers, but must do so without making users identifiable and without releasing the shared information for secondary uses or combination with data acquired elsewhere. Users who desire it should be offered the opportunity to contact advertisers without disclosing their profiles.

The fourth special case involves linkages between personal information and spatial management, including both generalized surveillance and processes for authentication of access. A regime of discontinuity-based privacy protection requires legal, policy, and technical interventions aimed at preserving adequate spatial privacy for situated subjects. Recall from Chapter 6 that the spatial-privacy interest operates in public spaces as well as private ones. In a networked information society, protection for spatial privacy requires strict limits on the retention of data establishing presence in most public spaces and in many technically private spaces that serve public functions (for example, a privately owned shopping center or the student commons at a private university). Transfers of such data during the term of its retention would be subject to the basic data-fiduciary model, and data-fiduciary rules also would operate to limit the real-time correlation of access records with stored profile data gathered from other sources. Both on- and off-line, surveillance should be visible; the persons or entities conducting it should be identified with particularity; and the rules governing the retention and processing of surveillance data should be publicly disclosed.

Finally, each of these proposed regimes also intersects with the problem of government data collection and use for law enforcement and national security purposes. As a practical matter, any privately held data set is potentially subject to compelled production, and we have seen that many government entities also participate actively in markets for personal information. A system of data-destruction mandates would not eliminate the latter activity, but instead would simply give the government incentive to acquire data before its destruction. Government practices with respect to personal information span a vast spectrum, and this chapter is already long. Many government uses of personal information are not different in kind from commercial uses and could be subjected to similar privacy rules. For present purposes, it is sufficient to note that

the most pressing government needs for access to accurate personal information in real time, which relate to law enforcement and counterterrorism surveillance, are not inconsistent with the principled development and application of a just-aggregation principle. It is just such a principle that the U.S. system of constitutional and statutory protections has attempted to achieve by erecting procedural barriers that revolve around particularized showing of a need for access. In the post-9/11 world, many of those protections have eroded. Lawmakers, enforcement officials, judges, and public opinion have become increasingly willing to accept the argument that every piece of information, however seemingly innocuous, may reveal a threat to public safety when placed in context. Yet that proposition is too often asserted rather than argued for.

The push toward more complete profiling in the interest of security reflects a particular philosophy of risk management, which holds that risk is most usefully conceptualized as an inverse function of logical completeness in information systems. But the relationship between information processing and risk is much more complicated than that view acknowledges. Events in the post-9/11 world reveal a dialectical relationship between new technological methods of managing risks and risks that new technological methods create. Large-scale data mining and complex, automated systems for managing critical infrastructures rely heavily on algorithms that align and systematize the meanings of data about people and events. Formally, such systems approximate the requirement of logical completeness, an approximation that becomes stronger as more and more data are collected. Much evidence suggests, however, that relying on such techniques to the exclusion of human judgment does not eliminate the risk of system failure, but instead magnifies the probability that system failures will be large and catastrophic. The U.S. government's development of a profile-based system for screening airline passengers inspired the "Carnival Booth" study, in which a pair of MIT-based researchers demonstrated how a terrorist group might defeat the screening system by hiding its agents within designated low-risk groups.³⁹ In 2009, the Washington, D.C., Metro system's exclusive reliance on an automated network-management system produced the deadliest subway crash in U.S. history.⁴⁰ Security experts believe that many other critical infrastructure systems are vulnerable to similar disruptions.

Chapter 8 explored some of the ways in which discourses of secrecy and spectacle underwrite our information policy; automated information processes intersect with those discourses in ways that can increase the risks of harm. Regimes of secrecy premised on need-to-know access to critical information increase the likelihood of groupthink and reduce the likelihood that critical perspectives will be brought to bear on security practices. The Carnival Booth study demonstrates powerfully that insider bias may reinforce the shortcomings of automated systems rather than correct for them. Meanwhile, the events leading up to the attempted bombing of a Detroit-bound airliner in December 2009 demonstrate that the U.S. security apparatus retains an astonishing capacity to ignore the results of human intelligence gathering.⁴¹ At the same time, the powerful conceptual link between comprehensive, rationalized information processing and security feeds the public demand for visible, information-intensive countermeasures without regard to whether they are also the most effective. As a result, security processes may tend to emphasize visibility over efficacy, or what Bruce Schneier has called "security theater."⁴²

Security planning to minimize the likelihood of catastrophic harm requires due regard for the risks of too much information, too much automation, and too much secrecy. Many security experts argue that humans are the best threat detectors because the insistently analog human brain draws connections that automated algorithms may not. The surprising level of consensus on that view suggests that the right response to contemporary security threats may be counterintuitive: less reliance on predictive profiling and more emphasis on heterogeneous and often redundant layers of protection. For present purposes, it seems sound to conclude that the regime of privacy protection sketched in this chapter would not make us less secure, and might produce the opposite effect.

Legal and technical privacy rules animated by a just-aggregation principle would work to produce a networked information society characterized by respect for what Helen Nissenbaum calls context-relative informational norms.⁴³ Critically, this is so whether or not those rules are perfectly enforceable. Here, privacy theory and policy can draw useful lessons from the copyright experience. Legal prohibitions on infringement are relatively ineffective at preventing P2P file sharing of copyrighted sound recordings, but robust markets for digital music nonetheless have emerged, and the widespread availability of lawful, affordable access has supported the development of norms favoring payment. A system of privacy laws will always remain vulnerable to abuse. But the processes of norm formation do not run only one way; privacy expectations are shaped not only by what is possible, but also by discourses about the content of legal rights and the nature of good engineering practice. Meaningful privacy protection has been difficult to attain because we as a society have been unwilling to commit to it either formally or intellectually. If such a commitment could be made, there is every reason to think that, over time, a rigorous, principled commitment to just aggregation would generate its own supporting discourses and norms. Just as in the case of semantically discontinuous copyright, moreover, a privacy regime founded on principles of just aggregation likely would enhance the legitimacy of the surveillance practices that need to remain in place.

Material Practice within Technical Landscapes

Chapter 8 argued that the play of everyday practice has an important material dimension that requires room for experimentation and play by situated users of networked information technologies. Emergent regimes of authorization seek to stabilize commercial relationships and public functions in a way that systematically minimizes breathing room for everyday practice, and that threatens important social values. Rather than automatically reinforcing such regimes, laws governing copyright, trade secrecy, and privacy must work together to balance fixity and play. To promote semantic discontinuity, legal and technical rules governing interconnection should seek to foster a heterogeneous, imperfect technical landscape that allows scope for the play of everyday material practice while maintaining protection for the privacy of situated users.

In technical and policy discussions about the design of network architectures, the interplay between logical skepticism and technocratic confidence revolves around the tension between architectures that are “open,” in the sense that no central decision maker controls their interoperability with networks, platforms, and tools, and architectures that are “closed.” Technical communities and policy communities experience that tension differently. However, dynamics

within each community contribute importantly to the emergence and gradual entrenchment of regimes of authorization.

Computer scientists and technology designers are inclined to view technical barriers to interoperability as artificial constraints to be overcome. That conviction derives partly from the information-processing imperative, already discussed. It also derives from a commitment to seamless, interoperable design that is both intellectual and aesthetic, and that is deeply internalized in the technoculture of computer science and engineering. Seamless interoperability and uninterrupted semantic flow are central goals in the theory and practice of network design. Belief in the fundamentally artificial nature of barriers to data interchange coexists, sometimes uneasily, with technocratic confidence in the possibility of defining increasingly granular, code-based rules for authorizing flows of information. As we have seen, that confidence manifests across each of the domains that this book has explored, in the development of rights-specification languages, surveillance systems, ubiquitous-computing systems, search algorithms, and so on. For technologists, the commitments to the foundational importance of openness and to the tantalizing possibility of control are reconcilable within a normative framework that expects, and indeed demands, continual challenge to reigning theoretical and technical frameworks. Most also recognize a role for legal and ethical rules that distinguish between productive inquiry and destructive vandalism.

For lawyers and policy makers, the considerations surrounding the interplay between openness and closure are more complex, reflecting the influence of additional, competing normative considerations that relate to social policy. Many (though not all) policy makers think that, other things being equal, open access to technical protocols promotes both innovation and competition. Other things often are not equal, however. As we saw in Chapter 8, the law protects technical secrets for a variety of reasons. In other cases, patent policy may support the development and licensing of technical protocols on a proprietary basis. In contemporary technology-policy debates, the intellectual property system's institutional support for closed systems based on proprietary technologies derives normative reinforcement from the seductive possibility of attaining more accurate regulation of behavior. Like Justice Holmes, we are skeptical of relying on the insubstantial reeds of virtue and internalized communal obligation to enforce rules of good conduct; unlike Justice Holmes, we are inclined to view the reed of law as equally insubstantial if there are technical measures that can accomplish the desired result. And once having committed to the importance of such systems, policy makers are inclined to think that they should not lightly be set aside. In particular, to the lawyers and businesspeople who play an instrumental role in defining and extending regimes of authorization, a professional culture that encourages the hacking of authorization systems seems exotic and alien.

Within the A2K paradigm, the ensuing controversies about access to and legal reinforcement of proprietary systems are most easily understood as debates about the relative merits of openness and closure. A2K advocates keenly appreciate the ways in which restrictions on technical accessibility have worked to tilt the playing field to the advantage of the economically and politically powerful. That history makes them enormously wary of legal involvement in standards setting, which they view as vulnerable to political capture, and of

mandated restrictions on the technical accessibility of information systems, which they view as inevitably disadvantaging the powerless. Scholars like Jack Balkin, Yochai Benkler, and James Boyle have argued persuasively that in the realm of intellectual property, open access to networks, information, and technical protocols promotes not only innovation and competition, but also important equality-related goals, including freedom of expression and access to the fruits of technical innovation. Most A2K advocates acknowledge that unauthorized access to closed systems may legitimately be prohibited in some circumstances, but they would drastically narrow the law's protection of publicly available platforms for copyrightable content and other information services. They have advocated the adoption of national technology policies mandating technical openness in certain core capabilities, and they are inclined to view most legal restrictions on interconnection as normatively unjustifiable. When confronted with technical mash-ups that recast personal information about network users in new ways—for example, merging the “Twitter stream” with global positioning data to pinpoint users precisely in space and time—they have been inclined to praise the technical creativity involved and to overlook or excuse the privacy implications.

If we interrogate this binary framing of the relationship between openness and closure from the perspective of everyday practice, things become more complicated. The play of everyday practice thrives when openness and closure are in balance. Emerging regimes of authorization threaten the play of everyday practice not because they implement universal closure, but more precisely because of the ways that they change the patterns of openness and closure that everyday practice requires to thrive. Regimes of authorization establish closed circuits of information flow governed by their own internal logics. The patterns of information flow created by copyright management and security protocols produce important and highly artificial discontinuities for network users, subjecting them to technical and transactional barriers that interfere with creative and material practice. Yet regimes of authorization also benefit from openness with respect to the collection and flow of personal information about users and user communities. Openness and closure together supply the foundation for the dynamics of transparency and exposure discussed in Chapter 6, and eliminating only legal protection for closure would not rectify the problems that openness creates.

Consider Google's recent entry in the social-networking field, Google Buzz, which trumpeted its adoption of open-data standards. The most controversial aspect of Google Buzz was Google's decision to combine data streams from its Buzz and Gmail products and to display users' top Gmail correspondents as their publicly disclosed Buzz “friends.” In response to public outrage over this unexpected blending of private and public, Google gave users a way to opt out of the default settings.⁴⁴ Substantive privacy protections such as those described in the previous section would limit Google's ability unilaterally to make such disclosures. But the open-data architecture of Google Buzz was no different from the proprietary architecture of market leader Facebook in one critical respect: it sharply limited users' power to create, maintain, and revise privacy-protective boundaries in context-specific ways. The Twitter stream example described above similarly exploits user powerlessness, a point powerfully demonstrated by the hack of Foursquare's constant stream of updates on

users' whereabouts to generate Please Rob Me, a site that linked users' out-and-about updates with their home cities.⁴⁵

When confronted with the privacy problems that unrestricted technical openness can create, A2K advocates tend to become sudden and unaccountable believers in the market's invisible hand. They argue that if platforms that allow social networking while protecting personal information are so desirable, users will create them, empowering new online communities to which other users will flock. As we saw in Part IV, that argument is structurally naive. The social and political effects of logical openness must be assessed in the context of a market structure that rewards transparency-promoting interconnection. Even new technical offerings touted as empowering users tend to harden their positions on transparency and exposure as they migrate out of the start-up phase. In addition, as James Grimmelmann explains, "The design of social network sites plays into plenty of well-understood social cognitive biases" by "activat[ing] the subconscious cues that make users think they are interacting within bounded, closed, private spaces."⁴⁶ Under current legal and market conditions, more effective tools for managing personal boundaries online are disfavored.

The A2K narrative about openness also tends to overlook the fact that even communities of practice organized around principles of technical openness and seamless interconnection sometimes pursue other values. A central tenet of open-source coding practice is that when serious disagreements about project direction arise, standards for the project can be "forked." Within the official discourse of open-source software-engineering practice, the possibility of forking serves as an important meritocratic corrective to the path-dependent engineering process. Yet forking also has other, geographic and political implications that are less well explored. The choice to fork an evolving protocol might be desirable precisely because it offers a choice to enable local platforms tailored to situated users' particular needs. Nor have open-source software designers fully rejected the closed systems characteristic of regimes of copyright authorization. In the DeCSS litigation, defendants argued that circumvention of the copy-protection system for DVDs was intended to create an open-source DVD player, and was necessary because no such player was available. Obtaining licensed access to the CSS technology would have required an agreement to embed robust, secret functionality at the core of an otherwise open product. While the two approaches are not incompatible from a technical perspective, some had raised larger questions about their philosophical compatibility. By the time the DeCSS case went to trial, however, the DVD Copy Control Association had granted two such licenses, evidence that at least two groups of developers had found the conflict to be reconcilable.⁴⁷

These examples in turn suggest that the equation of logical openness with political freedom is too simple. Situated users value openness very highly, but many can neither fully embrace standardization nor abandon dialogue with closure. As we saw in Chapter 7, moreover, the alignment of unrestricted technical openness with political and expressive freedom simply restates the traditional liberal preoccupation with liberty and constraint. The play of everyday practice requires no such perfect alignment. Play is not, and could not be, wholly liberated from circumstantial constraint; it follows, then, that circumstantial constraints need not foreclose meaningful opportunities for play. The question is not whether constraints should exist at all, but how to locate them in

a way that most effectively promotes all aspects of human flourishing. Wherever they are located, they will be challenged, but that does not necessarily make all constraints illegitimate. (Sometimes, transgression is just transgression.)

It is useful to disaggregate the problem of openness and closure into two narrower, functional categories subsumed within the paradigm of unauthorized access. One category of decisions concerns whether and when to penalize the act of unauthorized access to a closed system. The other concerns whether and when to enforce coordination around a closed standard. For regimes of authorization to succeed, the law must support both types of action. Many lawsuits that are framed as involving unauthorized access in fact involve the intersection of the “coordinated standardization” and “unauthorized access” categories, and those lawsuits suggest a template for resolution of the regulatory dilemma that the push for technical openness creates. The complex interrelationship between everyday practice and technical accessibility requires a regulatory landscape designed both to encourage certain kinds of interconnection and to promote certain kinds of closure.

Consider the recent litigation involving the RealNetworks RealDVD media player. RealNetworks designed a system that would enable users to play DRM-protected prerecorded DVDs, but it did not permit wholly unrestricted access to the content. Instead, it sought to provide access that, while it neither conformed in all respects to the applicable proprietary standard nor was authorized by the DVD Copy Control Association, nonetheless would provide meaningful copyright protection.⁴⁸ Plaintiffs objected not because the media player created an increased risk of infringement—by any objective standard, it did not—but rather because its development and distribution flouted the dominance of their regime of authorization. Although the suit was framed as one seeking to enjoin the provision of circumvention tools, the real dispute concerned the extent of legal support for privatized standardization—the extent to which law should delegate irrevocably to private actors the authority to specify how much content protection is enough. RealNetworks lost because the court read the statutory delegation as absolute, but the law could approach questions of privatized standardization differently.

The critical underpinning of regimes of authorization is legally sanctioned coordination around a closed standard. It is such coordination that most directly threatens human flourishing in the networked information society, and that a justice-promoting information policy should seek to neutralize. But the RealNetworks dispute suggests that the law could allow interconnection while imposing other conditions on it. Situated users should enjoy broad freedom to repurpose networked digital artifacts, but that freedom should end where legitimate interests—in copyright, in national security, or in meaningful privacy—begin. To provide meaningful shelter for the play of everyday practice, the law also should seek to counter the hardening of regimes of authorization more directly, by defining baseline implementation standards designed to preserve interstitial complexity in the technical environment.

In the case of technical protections for copyright, policy makers should seek to develop legal rules that differentiate more conscientiously between modes of unauthorized access that promote true piracy and modes that further the play of everyday practice. First, legal prohibitions on the act of unauthor-

ized access should distinguish between circumvention for willful infringement and circumvention for the expanded set of lawful uses described earlier in this chapter. Second and correspondingly, legal prohibitions on the development and provision of circumvention tools should be narrowed to permit interconnection by new, unlicensed content services and media players that afford an adequate amount of protection against unrestricted reproduction and retransmission. Such prohibitions also should exempt the provision of circumvention tools designed simply to assist users in making lawful uses of technically protected content.

In addition, the law should decline to enforce copyright-protection regimes that unduly burden the play of everyday practice. Many current approaches to the design of copyright-protection systems attempt to satisfy users' desire for portability of media content without affording parallel flexibility to copy and remix. Emerging regimes of authorization encourage technology intermediaries to comply with technical and contractual restrictions to minimize their own exposure to liability. The law governing copyright protection systems should seek to reverse this polarity, giving both copyright owners and intermediaries incentives to design and implement systems that incorporate more tolerance for play. To claim the benefits of anticircumvention protection, copyright owners should be required to produce evidence of such design efforts. To claim the benefits of safe harbors from indirect-infringement liability, intermediaries should be required to show that their systems do not unduly restrict lawful uses. Such burdens are not unrealistic. Efforts by researchers and open-access advocates have shown that it is possible, for example, to define filtering protocols for user-generated content more or less restrictively.⁴⁹ The "least cost avoider" rationale for defining and enforcing intermediary obligations tends to magnify the importance of legal violations by end users; the more general point, which tends to get lost in enforcement-oriented discussions, is that cost considerations may make intermediaries an appropriate focus of regulatory leverage in either direction.

In the case of privacy, policy makers should develop regulatory interventions that differentiate between interconnection practices that magnify the transparency and exposure effects experienced by situated subjects, and other practices that offset or minimize such effects. First, the law should permit circumvention of technical-protection systems for proprietary social-networking and gaming platforms as necessary to enable users to make lawful use of their own information or transfer it to competing information platforms. Second, it should permit interconnection with proprietary platforms and services if and only if the new platforms and services enabled by the interconnection afford users adequate privacy protection. Such a rule would acknowledge the potential cultural and political value of technical mash-ups, but also require mash-up creators to introduce other protections to offset the new kinds of information that they make visible. For example, the creators of the Twitter stream/GPS mash-up described above might compensate for the increased geographic exposure that they create by limiting the accessibility of usernames so that only users specifically authorized to do so could connect a particular person to a particular location.

In addition, the law should discourage design decisions that unduly threaten the play of everyday practice by subjecting users to heightened trans-

parency and exposure. Here, policy makers can draw concrete lessons from more theoretical scholarship about the historically and contextually contingent trajectories of human-designed artifacts. The flattened categorical structures and the lack of context sensitivity that so-called social software routinely exhibits are not inevitable, but rather reflect both the circumstances under which social-networking platforms arose and the values of their designers and operators. Two circumstantial factors in particular are worth considering more carefully. First, the vast majority of early adopters of social-networking technologies were quite young. Many (though not all) such individuals operate within flatter social schema than adults do, and consequently have less experience managing the boundaries that separate contexts from one another. This does not mean that the young do not value privacy; researchers who study online youth culture have shown that teens and twentysomethings often care deeply about preserving the contextual integrity of their online disclosures.⁵⁰ Nor does it mean that we should all simply learn to practice selective amnesia toward embarrassing antics and disclosures, as some commentators have argued. It means, instead, that the online behaviors of those who are still learning to construct and manage personal boundaries should not supply the normative baseline for policy making.

Second, as danah boyd has observed, certain design features of popular social-networking platforms—their relatively rigid, algorithmic categorization of people, and their inability to facilitate certain kinds of contextual separation—likely reflect the predilections and dysfunctions of geek culture rather than the preferences of social-networking participants more generally.⁵¹ Whether to embrace those predilections is itself a choice, and one with large consequences. The failure to erect obstacles to the market-driven logics of transparency and exposure invites those logics to expand into the spaces where boundary management is impaired. The inability to reinforce contextual separation also intersects with and reinforces majority cultural norms; its “nothing to hide” ethos effectively privileges a way of being in the world that many people—immigrants bridging two cultures, gay and minority youth, people struggling to extricate themselves from difficult or abusive relationships—do not experience.

As in the case of copyright, regulators should seek to reverse the polarities in social-software markets that favor the provision of overly lax and contextually insensitive privacy-management features. In particular, regulators should pursue two kinds of intervention. First, they should require developers of social-networking services to implement strong pro-privacy default rules and to educate users on their importance. Second, they should promulgate standards regarding the substantive adequacy of privacy-management tools that create incentives to develop such tools and provide adequate instruction on their use. A conceptual template for this sort of regulation may be found in the movement for value-centered design, which stresses the iterative articulation of and engagement with normative values throughout the design process.⁵² By analogy to the doctrines that establish secondary liability for copyright infringement, one might imagine a rule establishing liability for providing privacy-management tools that do not enable a reasonable degree of contextual variation and that do not afford a reasonable level of control. What is reasonable would depend on the state of the art, but would be the subject of an obligation to make continuing improvements.

A combination of conditional interconnection privileges and value-driven design obligations would work to sustain and reinforce interstitial complexity in the networked information environment. This in turn would help preserve semantic discontinuity within networked physical and digital spaces, safeguarding the processes and practices through which culture moves and changes and through which embodied, situated subjectivity is formed. Such a regime would entail a type of constraint on innovation; as we have seen, however, innovation in the service of openness is not an unmitigated good. We accept without question that new drugs should be evaluated for their effects on human health; so too, new technologies should be evaluated for their effects on human flourishing. Judged according to that standard, the regime I have sketched fares well. Most minimally, it is preferable to both of the currently existing alternatives—to the constraints on innovation imposed by regimes of authorization, on the one hand, and to the constraints on evolving subjectivity that result from transparency and exposure, on the other. More fundamentally, it would focus the attention of policy makers and technologists on the important and difficult challenge of facilitating the play of everyday practice so that the situated subjects and communities who engage in it can thrive.

Notes

¹ See Nussbaum, *Frontiers of Justice*, 76-78.

² Ibid.

³ Shaver, “Defining and Measuring A2K,” 239.

⁴ See Tushnet, “I Put You There,” 921-24; 17 U.S.C. §1201(a)(1)(C).

⁵ For examples illustrating the range of producer responses to user innovation, see Fisher, “The Implications for Law of User Innovation,” 1435-41. On play rhetoric and user labor, see, for example, Banks & Humphreys, “The Labour of User Co-Creators”; Zwick, Bonsu, and Darmody, “Putting Users to Work.” For a similarly skeptical perspective on the culture of digital gaming, see Grimes & Feenberg, “Rationalizing Play.”

⁶ The automatic license for the preparation of sound recordings is set forth in 17 U.S.C. §115. For the automatic licensing provisions that apply to digital public performances, see 17 U.S.C. §114(d)(2) & (f). Section 114(f)(e)(2) prohibits collective rate negotiation for performances that do not qualify for the statutory license.

⁷ See Chon, “Intellectual Property ‘From Below’”; Okediji, “Sustainable Access to Copyrighted Digital Information Works.”

⁸ Chon, “Intellectual Property and the Development Divide” 2893.

⁹ Okediji, “Sustainable Access to Copyrighted Digital Information Works,” 180.

¹⁰ Berkman Center for Internet and Society, *Next Generation Connectivity*, 3.

¹¹ See Lee, “Warming Up to User Generated Content”; Wu, “Tolerated Use.”

¹² See, for example, Benkler, “Free as the Air to Common Use,” 361-63; Boyle, “The Second Enclosure Movement,” 58-69.

¹³ See Organization for Economic Cooperation and Development, “OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,” http://www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html; Directive 95/46/EC of the Europe-

an Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, 1995 O.J. (L 281) 31.

¹⁴ See Organization for Economic Cooperation and Development, “The 30th Anniversary of the OECD Privacy Guidelines,” http://www.oecd.org/document/35/0,3343,en_2649_34255_44488739_1_1_1_1,00.html; see also Organisation for Economic Co-operation and Development, *The Seoul Declaration for the Future of the Internet Economy*, 9-10, <http://www.oecd.org/dataoecd/49/28/40839436.pdf>.

¹⁵ See Citron, “Technological Due Process”; Citron, “Open Code Governance.”

¹⁶ See Pasquale, “Beyond Competition and Innovation”; see also Introna & Nissenbaum, “Shaping the Web.”

¹⁷ DeNardis, *Protocol Politics*, 219.

¹⁸ *Ibid.*, 225.

¹⁹ Holmes, “The Path of the Law,” 459-61.

²⁰ See Gibson, “Risk Aversion and Rights Accretion,” 887-906.

²¹ Lessig, *Free Culture*, 187.

²² See Landes & Posner, “Indefinitely Renewable Copyright”; Sprigman, “Reform(aliz)ing Copyright,” 553-56.

²³ See Elkin-Koren, “What Contracts Cannot Do,” 397-419.

²⁴ See, for example, *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 590-94 (1994) (“Although 2 Live Crew submitted uncontroverted affidavits on the question of market harm to the original, neither they, nor Acuff-Rose, introduced evidence or affidavits addressing the likely effect of 2 Live Crew’s parodic rap song on the market for a nonparody, rap version of ‘Oh, Pretty Woman.’”); *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1168 (9th Cir. 2007) (“[T]he district court did not make a finding that Google users have downloaded thumbnail images for cell phone use. This potential harm to Perfect 10’s market remains hypothetical.”).

²⁵ See Cohen, “*Lochner* in Cyberspace,” 504-14; Gordon, “On Owning Information,” 166-80.

²⁶ See Horwitz, *The Transformation of American Law*, 208-12.

²⁷ Nissenbaum, *Privacy in Context*, 127-57. For Litman’s argument about shelter for personal use of copyrighted works, see Litman, “Lawful Personal Use.” For the statutory definition of the public performance right, see 17 U.S.C. §101 (“To perform or display a work ‘publicly’ means—(1) to perform or display it at a place open to the public or at any place where a substantial number of persons outside of a normal circle of a family and its social acquaintances are gathered.”).

²⁸ See Cohen, “Creativity and Culture,” 1198-1205; Cohen, “Copyright, Commodification, and Culture,” 160-64.

²⁹ Compare Gramm-Leach-Bliley Act, Pub. L. No. 106–102, 113 Stat. 1437 (codified at 15 U.S.C. §6802), with Financial Services Act of 1999, H.R. No. 10.RH, §501 (version reported in the House).

³⁰ See Ohm, “Broken Promises of Privacy,” 1719-20, 1723-27; Latanya Sweeney, “Computational Disclosure Control: A Primer on Data Privacy Protection,” Jan. 8, 2001, <http://groups.csail.mit.edu/mac/classes/6.805/articles/privacy/sweeney-thesis-draft.pdf>. For the HIPAA deidentification regulations, see 45 C.F.R. §164.514.

³¹ See 45 C.F.R. §164.512(i); Institute of Medicine, *Beyond the HIPAA Privacy Rule: Enhancing Privacy, Improving Health Research*, 162-87, Washington, D.C.: National Academics Press, 2009.

³² An early and well-known statement of this view is Lessig, *Code*, 156-63. For a devastating rejoinder, see Rotenberg, “Fair Information Practices and the Architecture of Privacy.” See also Schwartz, “Beyond Lessig’s *Code* for Internet Privacy.”

- ³³ The classic study is Edwin Black, *IBM and the Holocaust*.
- ³⁴ See, for example, Schauer, *Profiles, Probabilities, and Stereotypes*, 199-223.
- ³⁵ One recent thought experiment that moves in this direction is Hoofnagle & Solove, "A Model Regime of Privacy Protection."
- ³⁶ Ohm, "Broken Promises of Privacy," 1765-71.
- ³⁷ See Acquisti & Gross, "Predicting Social Security Numbers from Public Data."
- ³⁸ I thank Allan Friedman for enlightening me on this point.
- ³⁹ See Robert O'Harrow, Jr., "Intricate Screening of Fliers in Works: Database Raises Privacy Concerns." *Washington Post*, Feb. 1, 2002; Chakrabarti & Strauss, "Carnival Booth."
- ⁴⁰ See Lena H. Sun & Lyndsey Layton, "At Least 6 Killed in Red Line Crash," *Washington Post*, June 23, 2009.
- ⁴¹ See Mark Mazzetti & Eric Lipton, "U.S. Spy Agencies Failed to Collate Clues on Terror," *New York Times*, Dec. 31, 2009; Scott Shane, "Wide U.S. Failures Helped Airliner Plot," *New York Times*, May 19, 2010.
- ⁴² Schneier, *Beyond Fear*, 38-40.
- ⁴³ See Nissenbaum, *Privacy in Context*, 140-48.
- ⁴⁴ See Miguel Helft, "Anger Leads to Apology from Google about Buzz," *New York Times*, Feb. 15, 2010.
- ⁴⁵ See Ryan Kim, "'Rob Me' Site Points Out Security Concerns," *San Francisco Chronicle*, Feb. 18, 2010.
- ⁴⁶ Grimmelmann, "Privacy as Product Safety," 803.
- ⁴⁷ See *Universal City Studios, Inc. v. Reimerdes*, 111 F. Supp. 2d 294, 311, 319-20 (S.D.N.Y. 2000), *affirmed*, *Universal City Studios, Inc. v. Corley*, 273 F.3d 429 (2d Cir. 2001).
- ⁴⁸ See *RealNetworks, Inc. v. DVD Copy Control Association*, 641 F. Supp. 2d 913, 924-27 (N.D. Cal. 2009).
- ⁴⁹ See, for example, Electronic Frontier Foundation et al., "Fair Use Principles for User Generated Content," http://www.eff.org/files/UGC_Fair_Use_Best_Practices_0.pdf. For a comprehensive discussion and evaluation of a variety of less restrictive designs for technical-protection systems, see Armstrong, "Digital Rights Management and the Process of Fair Use." See also Burk & Cohen, "Fair Use Infrastructure for Rights Management Systems."
- ⁵⁰ See Jones et al., "Everyday Life Online"; Livingstone, "Taking Risky Opportunities," 404-07; Moscardelli & Divine, "Adolescents' Concern for Privacy"; "Steeves & Webster, Closing the Barn Door," 14-15.
- ⁵¹ See boyd, "Autistic Social Software."
- ⁵² See, for example, the essays and case studies collected in Friedman, *Human Values and the Design of Computer Technology*.