October 1998

Chaos Prevailing on Every Continent: Towards a New Theory of Decentralized Decision-Making in Complex Systems

David G. Post

David R. Johnson

Follow this and additional works at: <https://scholarship.kentlaw.iit.edu/cklawreview>

Part of the [Law Commons](#)

Recommended Citation

David G. Post & David R. Johnson, *Chaos Prevailing on Every Continent: Towards a New Theory of Decentralized Decision-Making in Complex Systems*, 73 Chi.-Kent L. Rev. 1055 (1998).

Available at: <https://scholarship.kentlaw.iit.edu/cklawreview/vol73/iss4/5>

This Article is brought to you for free and open access by Scholarly Commons @ IIT Chicago-Kent College of Law. It has been accepted for inclusion in Chicago-Kent Law Review by an authorized editor of Scholarly Commons @ IIT Chicago-Kent College of Law. For more information, please contact jwenger@kentlaw.iit.edu, ebarney@kentlaw.iit.edu.

“CHAOS PREVAILING ON EVERY CONTINENT”: TOWARDS A NEW THEORY OF DECENTRALIZED DECISION-MAKING IN COMPLEX SYSTEMS*

DAVID G. POST** & DAVID R. JOHNSON***

INTRODUCTION

For the past several years, we have been focusing our attention on questions about the governance of cyberspace: How should the rules that apply to conduct on the global digital network be set, and who should be setting those rules? What polity or polities should we look to as a source of legitimate and welfare-enhancing rules for conduct there? Whose law governs the many terms of an Internet

* Presented at the Chicago-Kent College of Law/Cyberspace Law Institute Symposium on The Internet and Legal Theory, March 13-14, 1998. Thanks to Jeff Williams for his superb assistance in developing and implementing the complex systems simulations discussed here, to Joe LaBarge for research assistance, and to Jack Goldsmith, Avery Katz, Dan Klerman, Mark Lemley, Larry Lessig, Peter Menell, Dawn Nunziato, Steve Salop, Warren Schwartz, David Skeel, and Peter Swire for their comments on earlier drafts. We have presented versions of this paper at the Olin Law and Economics Symposium on International Economic Regulation at Georgetown University Law Center, April 5, 1997, the Aspen Institute's Annual Review: Internet as Paradigm conference June 19-20, 1997, and the Temple University Law School's faculty colloquium November 9, 1997, and we thank the many participants who have given us useful comments on those occasions. We also gratefully acknowledge our obvious debt to the work of Stuart Kauffman. Comments are welcomed.

Our title is taken from Wislawa Szymborska's poem *Psalm*, reprinted in *POEMS NEW AND COLLECTED 1957-1997*, at 148 (Stanislaw Barańczak & Clare Cavanagh trans., Harcourt Brace & Co. 1998).

Oh, the leaky boundaries of man-made states!
How many clouds float past them with impunity;
how much desert sand shifts from one land to another;
how many mountain pebbles tumble onto foreign soil
in provocative hops!

....

Oh, to register in detail, at a glance, the chaos
prevailing on every continent!

....

[H]ow can we talk of order overall
when the very placement of the stars
leaves us doubting just what shines for whom?

** Associate Professor of Law, Temple University Law School; Co-Director, Cyberspace Law Institute. dpost@vm.temple.edu.

*** Partner, Wilmer, Cutler & Pickering; Founder, Counsel Connect. david.johnson@counsel.com.

transaction—the amount of disclosure required, say, of someone offering goods or services for sale, the behaviors that constitute actionable “fraud,” the conditions under which a transaction can be rescinded, and the like? Which of the many plausibly applicable bodies of copyright law do we consult to determine whether a hyperlink on a World Wide Web page located on a server in France and constructed by a Filipino citizen, which points to a server in Brazil that contains material protected by German and French (but not Brazilian) copyright law, which is downloaded to a server in the United States and reposted to a Usenet newsgroup for worldwide distribution, constitutes a remediable infringement of copyright? To whose trademark law should we turn in order to resolve the question of whether registration of an Internet domain name, using a sequence of characters protected as a trademark under the law of some jurisdictions but not others, is a wrongful act subject to penalty?¹

It is no longer controversial to assert that these are difficult questions,² providing, at least, a wealth of wonderful exam

1. See, e.g., David R. Johnson & David G. Post, *Law and Borders—The Rise of Law in Cyberspace*, 48 STAN. L. REV. 1367 (1996) [hereinafter Johnson & Post, *Law and Borders*]; David R. Johnson & David G. Post, *The New “Civic Virtue” of the Internet*, in THE EMERGING INTERNET 23 (1998) [hereinafter Johnson & Post, *The New “Civic Virtue”*]; David G. Post, *Governing Cyberspace*, 43 WAYNE L. REV. 155 (1996); David G. Post, *The “Unsettled Paradox”: The Internet, the State, and the Consent of the Governed*, 5 IND. J. GLOBAL LEGAL STUD. 521 (1998) [hereinafter Post, *The “Unsettled Paradox”*]. These papers are all available on-line at <http://www.cli.org> and <http://www.temple.edu/lawschool/dpost.html>.

We have been (fairly) chided for being insufficiently precise in distinguishing between our descriptive and normative claims in earlier work. See Lawrence Lessig, *The Zones of Cyberspace*, 48 STAN. L. REV. 1403, 1407 (1996) (commenting on Johnson & Post, *Law and Borders*, *supra*, and suggesting that “[t]o argue that real space law should leave cyberspace alone one needs a normative argument—an argument about why it is good or right to leave cyberspace alone”); *id.* at 1411 (“[I]f the argument for deference [to the emergent rules of on-line communities themselves] that Johnson and Post here beg is a normative argument, we must say something more about the normative attractiveness of the world that cyberspace will be.”). Although the claim that we argued that “real space law should leave cyberspace alone” is somewhat overstated, Lessig’s main point is well-taken, and we offer the current paper in the spirit of constructive engagement with his critique.

2. There is no shortage of writing about the novel jurisdictional issues posed by Internet transactions. See, e.g., Richard H. Acker, *Choice-of-Law Questions in Cyberfraud*, 1996 U. CHI. LEGAL F. 437; Curtis A. Bradley, *Territorial Intellectual Property Rights in an Age of Globalism*, 37 VA. J. INT’L L. 505, 510-19 (1997); Dan L. Burk, *Federalism in Cyberspace*, 28 CONN. L. REV. 1095 (1996) [hereinafter Burk, *Federalism in Cyberspace*]; Dan L. Burk, *Patents in Cyberspace: Territoriality and Infringement on Global Computer Networks*, 68 TUL. L. REV. 1 (1993); Matthew R. Burnstein, *Conflicts on the Net: Choice of Law in Transnational Cyberspace*, 29 VAND. J. TRANSNAT’L L. 75 (1996); William S. Byassee, *Jurisdiction of Cyberspace: Applying Real World Precedent to the Virtual Community*, 30 WAKE FOREST L. REV. 197 (1995); James Alexander French & Rafael X. Zahraiddin, *The Difficulty of Enforcing Laws in the Extraterritorial Internet*, 1 NEXUS J. OPINION 99 (1996), available in LEXIS, Lawrev Library, Allrev File; Alexander Gigante, *Ice Patch on the Information Superhighway: Foreign Liability for Domestically Created Content*, 14 CARDOZO ARTS & ENT. L.J. 523 (1996); Jane C.

hypotheticals for courses in Civil Procedure, Conflict of Laws, or International Law. But we would assert (perhaps somewhat more controversially) that they are more fundamentally troublesome than that. Physical and geographical constraints—on the effectiveness of action and on the exercise of physical power, on the relationship between causes and effects and the intensity of connections of all kinds between individuals, on the ability to mark boundaries that delineate legally-significant actors from one another—are deeply rooted in our day-to-day experience in the world of atoms. Like other “laws of nature”—gravity, entropy, the conservation of mass and energy, and so on—they undergird our legal systems in a most fundamental way; we would, obviously, think very differently about any number of legal regimes and legal questions were we living in a gravityless or frictionless world.³ But we need pay little attention to these fundamental constraints in order to analyze and understand those regimes and questions; precisely because they are invariant across the entire legally-relevant universe they can be entirely ignored, and we may cheerfully leave questions about the design of legal systems in a world where mass and energy are *not* conserved for contemplation by writers of science fiction.

But the physical constraints of distance and geography are no longer invariant across the legal universe. One corner of that universe—cyberspace—is creating a realm of human interaction in

Ginsburg, *Copyright Without Borders? Choice of Forum and Choice of Law for Copyright Infringement in Cyberspace*, 15 CARDOZO ARTS & ENT. L.J. 153 (1997); David Nimmer, *Brains and Other Paraphernalia of the Digital Age*, 10 HARV. J.L. & TECH. 1 (1996); Henry H. Perritt, Jr., *Jurisdiction in Cyberspace*, 41 VILL. L. REV. 1 (1996); Brad A. Slutsky, *Jurisdiction over Commerce on the Internet* (last modified June 6, 1997) <<http://www.kslaw.com/menu/jurisdic.htm>>.

3. Professor Lessig has written extensively on the important (but generally unacknowledged) ways that “laws of nature” act as regulatory constraints and on the ways that cyberspace regulation will need to re-conceptualize the role of these natural laws. See, e.g., Lawrence Lessig, *The Constitution of Code: Limitations on Choice-Based Critiques of Cyberspace Regulation*, 5 COMM.LAW CONCEPTUS 181, 181 (1997) [hereinafter Lessig, *Constitution of Code*] (describing the three interacting forces through which social control can be achieved—law, social norms, and laws of nature: “That I can not see through walls is a constraint on my ability to snoop. That I can not read your mind is a constraint on my ability to know whether you are telling me the truth. That I can not lift large objects is a constraint on my ability to steal. Nature, in these ways, constrains behavior. Nature, in this sense, regulates.”); Lawrence Lessig, *Reading the Constitution in Cyberspace*, 45 EMORY L.J. 869, 875 (1996) (suggesting that the *absence* of physical constraints in cyberspace makes regulation more complex than in the world of atoms: “[T]he problem of cyberspace for constitutional law [is that] it leaves us without constraint enough; . . . we are, vis-à-vis the laws of nature in this new space, gods; and the problem with being gods is that we must choose.”); see also Lessig, *Constitution of Code*, *supra*, at 181 (noting that while most “regulation talk” focuses on law and, in recent years, social norms, “modern legal scholarship has not thought much about how nature regulates”).

which those constraints are disappearing entirely, in which physical location and physical space are becoming both indeterminate and functionally irrelevant. There, physical proximity is no longer a prime determinant of the relationship between cause and effect in human interaction, connections of all kinds between individuals are no longer formed by location-dependent processes, and boundaries drawn in physical space can no longer be easily discerned (if they can be discerned at all).⁴

Does any of this raise truly *new* questions about the regulation of human behavior? Perhaps not—it is, after all, hardly clear that there *are* any such questions regarding human affairs. But an understanding of how interactions on the global network are to be governed, and how law-making sovereignty in regard to these interactions is to be defined, will, we believe, require some rather creative thinking about a problem—the role played by these locational constraints in legal systems—that has been (largely) ignored until now.⁵

4. See, e.g., Burk, *Federalism in Cyberspace*, *supra* note 2, at 1098-99 (because the Internet protocols call for “geographically extended sharing of scattered resources,” users “may therefore be completely unaware where the resource being accessed is, in fact, physically located. So insensitive is the network to geography, that it is frequently impossible to determine the physical location of a resource or user [because] [s]uch information is unimportant to the network’s function . . . and the network’s design thus makes little provision for geographic discernment.” (footnote omitted)); A. Michael Froomkin, *The Internet as a Source of Regulatory Arbitrage*, in *BORDERS IN CYBERSPACE* 129, 142-55 (Brian Kahin & Charles Nesson eds., 1997) (Internet users, “as long as they share a common language and a reasonably rapid connection,” will generally be “indifferent to the physical location of those with whom they communicate,” allowing them to engage in “regulatory arbitrage”); Johnson & Post, *Law and Borders*, *supra* note 1, at 1370-72 (describing the a-geographical nature of cyberspace); Joel R. Reidenberg, *Governing Networks and Rule-Making in Cyberspace*, in *BORDERS IN CYBERSPACE*, *supra*, at 84, 85-87 (describing the destruction of territorial and substantive borders in cyberspace).

This idea is beginning to percolate through the judiciary. See, e.g., *American Libraries Ass’n v. Pataki*, 969 F. Supp. 160, 168-69 (S.D.N.Y. 1997) (“The unique nature of the Internet highlights the likelihood that a single actor might be subject to haphazard, uncoordinated, and even outright inconsistent regulation by states that the actor never intended to reach and possibly was unaware were being accessed. Typically, states’ jurisdictional limits are related to geography; geography, however, is a virtually meaningless construct on the Internet.”); *Digital Equip. Corp. v. Altavista Tech., Inc.*, 960 F. Supp. 456, 462-63 (D. Ma. 1997) (“The Internet has no territorial boundaries. . . . Physical boundaries typically have framed legal boundaries, in effect creating signposts that warn that we will be required after crossing to abide by different rules. To impose traditional territorial concepts on the commercial uses of the Internet has dramatic implications, opening the Web user up to inconsistent regulations throughout fifty states, indeed, throughout the globe.” (citation omitted)).

5. As Richard Ford has written:

There is no self-conscious legal conception of political space. Most legal and political theory focuses almost exclusively on the relationship between individuals and the state. Judges, policymakers, and scholars analogize decentralized governments and associations either to individuals, when considered vis-à-vis centralized government, or

This is, to put it mildly, a big subject; first principles have an annoying tendency to require thinking at a fairly high level of abstraction. We have chosen to begin by posing a relatively simple problem-solving problem that we call the "Gardener's Dilemma": How can one find a "good," or perhaps even the "best," configuration of a complex system (in this case, a garden), a collection of individual elements (plants) whose behavior (however measured) is dependent upon the behavior of many others within the system?⁶ This problem is a close analogue of many problems in law and social policy, which often call for a comparison of alternate future configurations of complex social systems and a choice among them.⁷ The question "How do we solve problems like the Gardener's Dilemma?" is therefore of some general interest—as is the well-known observation that many problems of this sort are computationally intractable, incapable of true solution by any known methods.⁸ Legal theory would, we believe, be enriched by more focused consideration of this conundrum, by paying additional attention to the study of various algorithms derived from the study of "complex adaptive systems" that can successfully operate on problems of this kind. We describe one such algorithm—what Stuart Kauffman calls "patching," the division of a system into non-overlapping but coupled self-optimizing parts—that can, in certain circumstances, efficiently identify better configurations of complex systems (even if they can never be guaranteed to reach the optimal result).⁹ This patching algorithm bears a striking resemblance to well-known models of "competitive federalism," and we describe some experimental results that suggest that one feature of patched systems—the "congruence" between inter-individual spillover effects and the boundaries of these decision-

to the state, when considered vis-à-vis their own members, but consider the development, population and demarcation of space to be irrelevant. . . .

Legal boundaries are often ignored because they are imagined to be either the product of aggregated individual choices or the administratively necessary segmentation of centralized governmental power. . . .

....

. . . [W]e have not one, but two tacit conceptions of space—space as opaque and space as transparent. On the one hand, we often implicitly see political space as natural and fixed. On the other hand, and often at the same time, we see political space as irrelevant.

Richard Thompson Ford, *The Boundaries of Race: Political Geography in Legal Analysis*, 107 HARV. L. REV. 1841, 1857-59 (1994).

6. See *infra* section I.A.

7. See *infra* section I.B.

8. See *infra* section I.C.

9. See *infra* Part II.

making patches—is a critical determinant of the effectiveness of patching as a problem-solving device.¹⁰ We conclude with some speculation as to the implications of these findings for the problem of Internet “governance,” and for existing models of federalism and related governance structures.¹¹

I. THE ANALYSIS OF COMPLEX SYSTEMS

A. *The Gardener's Dilemma*

Imagine a garden consisting of many different plants of many different species and a gardener who seeks to maximize some variable over the garden as a whole—total yield, for example. The gardener faces a particular decision: whether, with respect to each individual plant, to prune it back or leave it unpruned. How can our gardener find the “best” combination of pruned and unpruned plants, the configuration that will produce the most luxuriant growth overall?

The garden, we assume, has the following general characteristics. First, individuals are *heterogeneous*; the relationship between an individual's state (pruned or unpruned) and its growth is different for each plant; for some plants, pruning will increase growth (by reducing the diversion of scarce nutrients, water, sunlight, etc., into unnecessary foliage), while for others the “shock” of pruning will cause them to grow less vigorously. This heterogeneity may be due to differences among species—asparagus plants may react differently to pruning, on average, than tomatoes—and to differences between individuals of the same species (because of differences in overall health or vigor or genetic constitution, for example).

Second, there are substantial *spillover effects* between and among the individual plants, by which we mean that each individual's growth can be affected—positively or negatively—by the condition and growth of other plants.¹² The condition of an individual's neighbors

10. *See id.*

11. *See infra* Part III.

12. The notion of a spillover effect is similar to the familiar concept of an “externality.” In its most common usage, an “externality” describes a spillover effect that has the additional characteristic that it is not the subject of a market transaction. *See generally* ANDREAS A. PAPANDREOU, EXTERNALITY AND INSTITUTIONS 13 (1994) (“A History of the Notion of Externality”) (summarizing the various usages of the term within economics); Harold Demsetz, *Towards a Theory of Property Rights*, 57 AM. ECON. REV. 347, 348 (1967) (“Externality is an ambiguous concept. . . . What converts a harmful or beneficial effect into an externality is that the cost of bringing the effect to bear on the decisions of one or more of the interacting persons is too high to make it worthwhile”); Walter P. Heller & David A. Starrett, *On the Nature of Externalities*, in THEORY AND MEASUREMENT OF ECONOMIC EXTERNALITIES 9, 10 (Steven

will help determine, for example, the amount of sunlight that is likely to penetrate through to any individual plant, the amount of nutrients likely to remain in the immediate proximity in the soil, and the like, all of which will in turn affect the growth of that individual.

Third, each plant's response to being in one state or another (pruned/unpruned) is *endogenously determined*, i.e., a function of the state of some number other plants. For example, a plant's response to being pruned may depend on whether it is in an area of high or low sunlight; in the former, it may grow more vigorously if unpruned (which will allow it to take better advantage of the available sunlight), while with less sunlight available it may do better if unnecessary foliage is pruned away (or vice versa in some cases). And whether an individual plant is in an area of high or low sunlight in turn depends on the state (pruned/unpruned) of its neighbors.

Let us state this problem a bit more formally. The garden is a system consisting of some number (N) of individual *elements*—individual plants. Each element can be in one of only two possible *states*—for purposes of our problem, pruned or unpruned.¹³ The garden can contain at any time some elements in one state and some in the other; we will call each particular combination of pruned and unpruned plants a different *configuration* of the system.¹⁴ Because the elements of this system can be in one of only two possible states, we can represent any system configuration by a string of 1s and 0s, where a "1" indicates that a particular element is in the first ("pruned") state, a "0" the opposite, or by a diagram like Figure 1.

A.Y. Lin ed., 1976) (an externality exists where "the private economy lacks sufficient incentives to create a potential market in some good and the nonexistence of this market results in losses in Pareto efficiency"). Our use of the term "spillover effect" corresponds to what Andreas Papandreou calls the "broadest view of externality": a "general interdependence or any and all effects of others' actions on an agent." PAPANDREOU, *supra*, at 61. To avoid further muddying these muddy waters, we think it best to use the term "spillovers" to describe these inter-element effects.

13. We could define a richer "state space," if we wished, with more than two possible states for each element, but that would merely complicate the discussion that follows without changing the qualitative conclusions we are drawing from analysis of these systems.

14. In any system with N elements, each of which can take one of S possible states, there are S^N different system configurations. A 10 plant garden, in which each plant can be in one of two states, can take on $2^{10}=1024$ different configurations. Each configuration can be thought of as defining a single point in N -dimensional space—the "system space"—the axes of which are defined by the possible states of each of the N elements.

FIGURE 1

Visual depiction of a 100-element system organized, for convenience, into a square lattice. Open circles denote elements that are in one state ("pruned"), closed circles denote elements in the other state ("unpruned"). We can also represent this configuration by a 100-element binary string *viz.* [0, 0, 1, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, ...].

Each element's contribution to the system variable we are seeking to maximize is a function of both its own state and, because of inter-individual spillover effects, the state of some number of other elements; that is, a change in one element's state (from 0 to 1, or unpruned to pruned, or vice versa) affects both its own contribution and the contribution of a number of other plants to the overall yield of the garden. Conversely, each plant's contribution to overall yield is a function of both its own state and the state of a number of other plants in the garden.

Each configuration of the system produces some value for the system variable that we are seeking to maximize, and it is helpful to visualize a graph plotting the value of this system variable (aggregate yield in our example) against each different configuration of individual states (pruned and unpruned) that the system can be in. Such a map of the way that this variable changes as the states of individual elements change produces a kind of "landscape," a multi-dimensional terrain that rises to "peaks" of high yield in certain configurations and descends into "valleys" of low yield for other configurations of the elements.¹⁵

15. This idea of a "yield landscape" for the garden derives from the concept of a "fitness landscape," introduced by the population geneticist Sewall Wright to describe the distribution of the biological fitness—reproductive success—over the space defined by all possible genotypes (combinations of genetic material) within an evolving population. See STUART A. KAUFFMAN, *THE ORIGINS OF ORDER: SELF-ORGANIZATION AND SELECTION IN EVOLUTION* 33-120 (1993).

[hereinafter KAUFFMAN, ORIGINS OF ORDER]. The process of evolutionary adaptation is conceived of as a kind of "hill climbing via fitter mutants toward some local or global optimum," the process by which the genotype wanders over this fitness landscape in genotype space. See *id.* at 33; see also STUART KAUFFMAN, AT HOME IN THE UNIVERSE 149-89 (1995) [hereinafter KAUFFMAN, AT HOME IN THE UNIVERSE] (describing the landscape metaphor in detail).

To visualize such a "yield landscape," consider a simple, 2-element system. Each element has some contribution to overall yield in each of the two possible states; for example, the yield function for element X may be:

Yield Function, element X	
State	Yield Contribution
0	0.211
1	0.879

And for element Y:

Yield Function, element Y	
State	Yield Contribution
0	0.505
1	0.951

The aggregate yield for the system as a whole is thus given by the following table:

System Configuration	Aggregate Yield
0,0	0.716 [0.211+0.505]
0,1	1.162 [0.211+0.951]
1,0	1.384 [0.879+0.505]
1,1	1.830 [0.879+0.951]

We can plot this relationship on a graph with N+1=3 axes, with the vertical axis representing aggregate yield and the state of each of the N elements along N other axes. Each of the four possible system configurations can be represented by a single point in this N+1 dimensional space, and the "yield landscape" consists of the "height" of the various points on the fitness axis:

Note that this system has a yield "peak" at (1,1) and a yield "valley" at (0,0).

Unfortunately, because the "yield landscape" for an N-element system requires N+1 dimensions, it is impossible to actually display the fitness landscape graph for any system with N>2; conceptually, however, the idea of a plotted landscape can be applied to systems of any dimensionality.

The gardener's dilemma, then, is to find a way to identify the system configuration—the combination of individual state settings for each of the plants—that produces the maximum yield of the garden as a whole, the highest point on the yield landscape. How can the gardener do so?

B. The Law and the Garden

Before addressing that question, we should pause to consider the relevance of the Gardener's Dilemma to the Internet or to legal theory (the subject matters of this symposium). Why might we be interested in the way that the Gardener's Dilemma is (or is not) solved?

The problem posed by the Gardener's Dilemma is a highly general one, frequently encountered in the law. To decide whether a particular activity should be encouraged, discouraged, ignored, or prohibited by a legally-enforceable rule of some kind, we often ask some variant of the following question: "How will an individual's (or group's) performance of some activity (change of state) affect some measure of well-being (overall yield) within the population to which that individual belongs (the garden)?"

Consider the following entirely typical example, drawn from a recent case in which a federal court was asked to decide whether copyright law prohibits or permits college and university professors to photocopy articles from books or journals for distribution in "coursepacks" to their classes without authorization from the author or copyright holder.¹⁶ The relevant statutory provisions can be easily identified: the production of coursepacks constitutes an activity—"reproduc[ing] . . . copyrighted work[s] in copies"¹⁷—that is squarely within the exclusive rights of copyright holders; at the same time, the statute excuses "the fair use of . . . copyrighted work[s]," defined to include such uses as "reproduction in copies . . . for purposes such as criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research."¹⁸ Does coursepack production fall into this category or not?

Given indeterminate statutory language¹⁹ and no direct

16. See *Princeton Univ. Press v. Michigan Document Serv., Inc.*, 99 F.3d 1381, 1383 (6th Cir. 1996) (en banc).

17. Copyright Act, 17 U.S.C. § 106(1) (1994).

18. *Id.* § 107.

19. The statutory language governing the "fair use" inquiry is notoriously—though hardly

precedent on point,²⁰ the court looked, appropriately enough, to the purposes of copyright law (and the fair use exception) more generally in making its decision: Which rule will, in the constitutional phrase, best "promote the Progress of Science and the useful Arts"?²¹ Which rule will best "stimulate artistic creativity for the general public good"?²² Will placing this activity outside the copyright holder's exclusive rights increase the production of works of authorship (by allowing for greater dissemination of existing works among scholars and students), or will declaring this activity to be within those exclusive rights stimulate increased production of such works (by providing additional financial incentive for producers and/or publishers)?

As it turned out, the court was deeply divided on this question. A majority of the court was of the latter view:

[T]he district court was not persuaded that the creation of new works of scholarship would be stimulated by depriving publishers of the revenue stream derived from the sale of permissions. Neither are we. On the contrary, it seems to us, the destruction of this revenue stream can only have a deleterious effect upon the incentive to publish academic writings.²³

To the dissenters, on the other hand, the Copyright Act's "laudable societal objectives . . . to stimulate production of . . . original works for the benefit of the whole nation" will be "thwarted" unless professors were permitted to engage freely in this activity:²⁴

The majority's strict reading of the fair use doctrine promises to hinder scholastic progress nationwide. By charging permission fees on this kind of job, publishers will pass on expenses to colleges and universities that will, of course, pass such fees on to students. Students may also be harmed if added expenses and delays cause professors to opt against creating such specialized anthologies for

uniquely—indeterminate; one sitting federal judge has called the four statutory fair use factors the "fuzzball factors." See Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEGAL F. 207, 208. Indeed, the fact that eight judges of the Sixth Circuit found that the challenged use in this case was *not* fair use notwithstanding express reference in the statutory provisions to "multiple copies for classroom use" testifies to just how indeterminate the language of the statute is.

20. The *Princeton University Press* case was a case of first impression in the Sixth Circuit. See 99 F.3d at 1394 (Merritt, J., dissenting).

21. U.S. CONST. art. I, § 8.

22. *Princeton Univ. Press*, 99 F.3d at 1391 (the "ultimate aim [of copyright law] is . . . to stimulate artistic creativity for the general public good. . . . When technological change has rendered its literal terms ambiguous, the Copyright Act must be construed in light of this basic purpose." (internal quotation marks and citations omitted)).

23. *Id.*

24. See *id.* at 1393 (Martin, J., dissenting).

their courses. Even if professors attempt to reproduce the benefits of such a customized education, the added textbook cost to students is likely to be prohibitive.²⁵

Our point is not to choose sides in this particular debate, but to point out how similar the question all were asking—how will a prohibition on uncompensated copying in this context affect the production of scholarly works, and thereby the general welfare?—is to the Gardener's Dilemma. Deciding whether or not copyright law should permit professors to engage in this conduct presupposes an ability to say *something* about the way in which different "configurations" (of copying and no copying) within this system affect both the production of creative works and aggregate social well-being, something about the shape of a "well-being landscape" that is defined over different configurations of copying and no copying. Performance (or non-performance) of this activity surely has heterogeneous effects on individual well-being. Inter-individual spillover effects are certainly nontrivial; a professor's decision to photocopy a particular article, for example, affects the well-being (again, however one might define it) of her students, of the author of the works involved,²⁶ of other authors and potential authors,²⁷ of the

25. *Id.* at 1393-94 (Martin, J., dissenting); *see also id.* at 1394-97 (Merritt, J., dissenting) (arguing that the "essence of copyright is the promotion of learning," which will not be fostered by the majority's reading); *id.* at 1397-1411 (Ryan, J., dissenting) (the "more reasonable presumption" is not that "the practice of excerpting some materials harms the authors' rightful market and secures a benefit only to the excerpters" but that "society benefits from the additional circulation of ideas in the educational setting," and "speculat[ing]" that the defendants' copying either has "no impact on [the] incentives to authors to create new works, and may even *provide* authors incentive to write, thereby advancing the progress of science and the arts").

26. The effect of copying by any individual professor on the author's well-being may be a complicated function of the context in which the copying takes place. If the author derives compensation from the sale of copies of the work in question—as generally assumed in most analyses of copyright principles, *see, e.g.,* Trotter Hardy, *Property (and Copyright) in Cyberspace*, 1996 U. CHI. LEGAL F. 217, 221-22 (discussing the incentives for production of intellectual property that are unaffected by widespread copying, such as the "incentive for university professors . . . to achieve recognition though widespread distribution of their scholarly work"); William M. Landes & Richard A. Posner, *An Economic Analysis of Copyright Law*, 17 J. LEGAL STUD. 325, 327 (1989) ("[T]he expected return [for a new work to be created is] typically, and we shall assume exclusively, [derived] from the sale of copies.")—coursepack copying may reduce the author's well-being. Many authors, however, and especially those in markets with significant reputational components, may derive compensation for creating new works in other ways, *e.g.,* the reputational benefits that may accrue from wide distribution of copies of their work. *See id.* at 331 (observing that many authors derive substantial benefits from publication over and above income derived from the sale of copies, *e.g.,* "a higher salary for a professor who publishes than for one who does not, or greater consulting income" (footnote omitted)); Jessica Litman, *Copyright Noncompliance (Or Why We Can't "Just Say Yes" to Licensing)*, 29 N.Y.U. J. INT'L L. & POL. 237, 246-51 (1997) (suggesting that the current array of incentives may be sufficient to generate high-quality creative works on the Internet notwithstanding the inability of many producers to charge for copies of their work). *See*

employees and stockholders of the publishing house that published the author's work, of the copy shop that receives the copying revenue, and so on. And the richness of these couplings and interconnections means that the effect of copying on any individual's well-being is at least partially endogenously determined, dependent upon the copying activities of others (i.e., the "state" of other "elements" of the system).²⁸ If, like the gardener, we are looking for those particular

generally David Friedman, *Standards as Intellectual Property: An Economic Approach*, 19 U. DAYTON L. REV. 1109, 1115-16 (1994) (discussing various factors that may influence the "supply curve" for producers of intellectual property); Peter S. Menell, *The Challenges of Reforming Intellectual Property Protection for Computer Software*, 94 COLUM. L. REV. 2644, 2647 (1994) (arguing that product marketing, support, and reputation can be significant forces to protect market share notwithstanding widespread copying); Tom G. Palmer, *Intellectual Property: A Non-Posnerian Law and Economics Approach*, 12 HAMLINE L. REV. 261, 287-302 (1989) (surveying the incentives that may exist for the production of intellectual property in the absence of prohibitions against copying).

27. A central tenet of traditional copyright theory, of course, is the assumption that the ability to copy without compensation has a negative impact on the incentives of future authors—that the professor's copying, in some small measure, adversely affects those considering the investment of time and energy into the production of similar works in the future.

28. There are innumerable ways in which the relative costs and benefits of copying or no copying for any individual professor are endogenously determined by the "system configuration" itself, i.e., the state (copy/no copy) of others. Benefits derived from any particular act of copying are reduced to the extent a professor's students themselves are in the "copy" state (i.e., have made copies of the works in question, thereby obviating the need for the coursepack), as well as to the extent that those students can easily and at low cost place themselves in that state by obtaining copies of the work(s) in question from others who are themselves in the "copy" state (e.g., from others who have placed the materials in an easily accessible place such as a World Wide Web site). Over time, the "configuration" of professors—the distribution of "copy/no copy" states among this group—may generate a number of different responses by authors and publishers that in turn affect the benefits that professors receive from changing their state with regard to copying. For example, if professors are generally in the "no copy" state (whether because of a copyright rule prohibiting professorial coursepack production or otherwise), that alters the incentives of authors seeking wider distribution of their work to students to find alternate distributional channels (thereby reducing the benefit each professor receives from copying still further). Alternatively, widespread professorial production of coursepacks may increase the supply of copying services in and around universities, which may have the effect of allowing students to produce their *own* copies of required articles at lower cost, or it may stimulate other institutional responses that alter the professorial cost-benefit calculus in other ways (such as collecting societies acting on the authors' or publishers' behalf). See, e.g., Robert P. Merges, *Contracting Into Liability Rules: Intellectual Property Rights and Collective Rights Organizations*, 84 CAL. L. REV. 1293 (1996) (describing emergence of collecting societies as a response to initially inefficient property rights entitlements). In other words, the configuration of "copy/no copy" across university professors defines the problem that *others* are trying to solve, and the responses of those others alter, in turn, the problem the professors are trying to solve. And finally, just to square the circle of endogenous effects, the configuration of "copy/no copy" across this system itself affects the likelihood that the copying will be deemed "fair use" (further altering the costs and benefits of any individual professor's copying activities) in a number of ways. First, the incidence of coursepack copying affects (a) the likelihood that any single instance of such copying can be detected, which affects in turn (b) both the (static) costs of copying facing any individual professor and the (dynamic) likelihood that any monitoring and licensing scheme will be undertaken by authors, publishers, or the state, which affects in turn (c) whether the copying

configurations of performance and nonperformance that are “better” than others from the perspective of the system as a whole, how can we find them?

This copyright illustration is hardly exceptional; while we would not suggest that *all* normative legal questions take this form, we do think that a large number of problems in legal policy share its essential features. It is of no small interest, then, to see if we can answer the question with which we began: How do we solve the Gardener’s Dilemma? How can we identify “better” or “worse” configurations of such systems? How do we (or the judges of the Sixth Circuit) decide whether professorial production of coursepacks will or will not stimulate scholarly creativity and thereby better promote the “Progress of Science and the useful Arts”?

C. Solving the Gardener’s Dilemma

Notwithstanding the fact that the Gardener’s Dilemma does not appear to be a particularly difficult problem, the manner in which problems of this kind are solved—or whether, indeed, they *are* solved—is not at all well understood. The intuition involved here is relatively straightforward. Imagine that you are trying to solve the Gardener’s Dilemma one plant at a time; that is, moving one plant at a time across the garden, you determine whether pruning the plant in question would increase or decrease the yield of the garden, you mark the plant accordingly, and then you move on to the next. The growth of the first plant you encounter is a function (by our spillover assumption) of the condition of some others; thus, solving the Problem of the First Plant—determining whether pruning will cause it to grow more, or less, luxuriantly—therefore requires knowledge about the state of those other plants. But the state of those others will depend on the state of the first plant; the decision you are making now (for plant 1) changes the conditions defining the problem that you will be trying to solve later (i.e., whether or not to prune plant X).

falls into the fair use category. *Compare Princeton Univ. Press*, 99 F.3d at 1387 n.4 (existence of an “established license fee system is highly relevant” to fair use determination), *and American Geophysical Union v. Texaco Inc.*, 60 F.3d 913, 931 (2d Cir. 1994) (“[I]t is sensible that a particular unauthorized use should be considered ‘more fair’ when there is no ready market or means to pay for the use, while such an unauthorized use should be considered ‘less fair’ when there is a ready market or means to pay for the use.”), *with Princeton Univ. Press*, 99 F.3d at 1407-09 (Ryan, J., dissenting) (criticizing the circular nature of this logic). And second, there is the additional consideration that if widespread coursepack production without compensation to authors is the norm, courts may be more likely to deem such copying “reasonable” and within the fair use exception than otherwise.

It seems impossible to solve one without solving the other. The problem has a kind of nightmarish recursiveness about it; each different configuration of the system—and there are an enormous number of such configurations²⁹—seems to present a different problem to be solved and determining the system-wide maximum seems to require solving all of them simultaneously. The system, in Stuart Kauffman's words, is "caught in a web of conflicting constraints" in which "each small part of the system affects other parts of the whole system," and "changing [the state of a single element] . . . will have effects that ripple throughout the system."³⁰ A small change in the state of some elements may, by virtue of these complex cross-couplings between elements, have large impacts on the yield of the system as a whole.³¹

29. See *supra* note 14. A small ($N=100$) system with two possible states for each element ($S=2$) can take any one of 2^{100} possible configurations—a number approaching in magnitude the number of elementary particles in the known universe.

30. KAUFFMAN, AT HOME IN THE UNIVERSE, *supra* note 15, at 173. Lon Fuller called these kinds of problems "polycentric" (following Michael Polanyi's use of that term in THE LOGIC OF LIBERTY: REFLECTIONS AND REJOINDERS (1951)) and offered the following example:

Some months ago a wealthy lady by the name of Timken died in New York leaving a valuable, but somewhat miscellaneous, collection of paintings to the Metropolitan Museum and the National Gallery "in equal shares," her will indicating no particular apportionment. When the will was probated the judge remarked something to the effect that the parties seemed to be confronted with a real problem. . . . What makes this problem of effecting an equal division of the paintings a polycentric task? *It lies in the fact that the disposition of any single painting has implications for the proper disposition of every other painting.* If it gets the Renoir, the Gallery may be less eager for the Cezanne but all the more eager for the Bellows, etc.

Lon L. Fuller, *The Forms and Limits of Adjudication*, 92 HARV. L. REV. 353, 394 (1978) (emphasis added) (footnote omitted) [hereinafter Fuller, *The Forms and Limits of Adjudication*]. Fuller gave other examples of polycentric problems of this kind, from setting prices and wages within a managed economy to produce a proper flow of goods, where "[a] rise in the price of aluminum may affect in varying degrees the demand for, and therefore the proper price of, thirty kinds of steel, twenty kinds of plastics, an infinitude of woods, other metals, etc." and where "[e]ach of these separate effects may have its own complex repercussions in the economy," *id.*, to redrawing the boundaries of election districts, assigning the players of a football team to their respective positions, and allocating air routes among cities. See Lon L. Fuller, *Adjudication and the Rule of Law*, 1960 PROC. AM. SOC'Y INT'L LAW 1, 3 [hereinafter Fuller, *Adjudication and the Rule of Law*]. As Fuller himself noted, these problems are ubiquitous in the law. See *id.* at 7 ("I hope I shall not appear to be overworking the concept of polycentricity if I say that all community is polycentric in nature, as indeed are all living relationships.").

We thank Professors Jeffrey Dunoff and Michael Libonati for pointing out this theme in Fuller's work to us.

31. Fitness landscapes in complex, interconnected systems will, in other words, generally be extremely *rugged*; changing the state of a single element in such a system—taking a single "step" in the multi-dimensional system space—may affect the fitness of many (or, at the limit, all) other system elements, and can, as a result, produce dramatic changes in aggregate system fitness. Thus, aggregate system fitness may be highly sensitive to small changes in the states of the individual elements; points on the landscape that are "close together," i.e., separated in

You might think, as you laboriously make your way across the garden, that there must be sophisticated mathematical techniques that can be brought to bear on this problem (and you might take some comfort from knowing that such techniques exist, even if you are not yourself competent to utilize them). But, interestingly, that is not the case. There is a rich literature, primarily (but not exclusively) in the fields of computer science and economics, demonstrating that a wide range of problems characterized by the features we have given to our garden—heterogeneous individual elements, endogenously determined individual responses to changes in state, and inter-element spillover effects of substantial magnitude—are “computationally intractable”; they are effectively insoluble by any known techniques by virtue of the fact that the number of discrete computational steps required to compute a solution increases *exponentially* with the “size” of the problem, the “dimensionality” of the system space over which the problem is defined.³² In the “worst case”—one in which the interconnections among elements are sufficiently rich and complex such that each element’s state affects the growth of *all* other elements in the system and *vice versa*—we may be able to solve for the highest point on the system’s “yield landscape” only by searching over the *entire* configuration space; by the simple but inexorable logic of exponential scaling, this is impossible for even relatively small systems, for the number of possible configurations becomes unimaginably vast very quickly as *N* increases.³³

What we have, in any complex interconnected system, is a genuinely hard problem: the system can take any one of an astronomical number of different configurations and may, via the

multi-dimensional space by only one or a small number of state changes, may have very different fitness values, producing a jagged shape to the fitness landscape. See generally KAUFFMAN, *AT HOME IN THE UNIVERSE*, *supra* note 15, at 173-206 (a non-technical description of the relationship between conflicting constraints in interconnected systems and the ruggedness of their fitness landscapes); KAUFFMAN, *ORIGINS OF ORDER*, *supra* note 15, at 33-120 (a more technical discussion of the structure and biological implications of rugged fitness landscapes).

32. Professor Rust provides an excellent overview of the development of “impossibility theorems” in these various disciplines. See John Rust, *Dealing with the Complexity of Economic Calculations*, in *FUNDAMENTAL LIMITS TO KNOWLEDGE IN ECONOMICS* (J.F. Traub & S.N. Durlauf eds., forthcoming 1998), available on-line (last modified Oct. 17, 1997) <<http://econwpa.wustl.edu/eprints/comp/papers/9610/9610002.abs>>. Rust defines “computationally intractable” problems as those “for which the lower bound on the computation cost increases exponentially with the problem dimension.” *Id.* Examples of problems exhibiting this “curse of dimensionality” include the “traveling salesman” problem familiar to students of operations research, computation of equilibrium solutions to multi-person games, and calculation of the fundamental “Walrasian auction” equilibrium in microeconomics. See *id.*

33. See *supra* note 29.

"ripple effects" of changes throughout the system, respond to small changes of the state of those elements with large changes in whatever system-wide variable we seek to measure. We are looking for the highest point in the landscape defined over such a system, but can be assured of finding it only if we search the entire state space, an impossible task.³⁴ This "curse of dimensionality"³⁵ is not some odd bit of mathematical trivia;³⁶ it represents a substantial—and quite possibly a fundamental—constraint on our ability to solve a wide

34. This is not, we emphasize, because we do not have enough information about the relationships between the elements or about the manner in which an individual's state affects its own utility function and the utility of others, nor is it because the values of the relevant variables are somehow stochastically determined. Even when the Gardener's Dilemma is fully specified and deterministic—when we know, in advance, all determinants of each element's fitness and the way that each individual plant will react to being pruned or being left unpruned, and the way that each individual plant would be affected were the state of its neighbors to change—our problem with respect to finding the configuration that produces optimal yield for the system as a whole remains. The indeterminacy is a function of the inherent complexity of the system and the existence of conflicting constraints among the individual elements.

35. See Rust, *supra* note 32.

36. The impossibility of computing "equilibrium solutions" to the problems posed by any modern economic system was central to Friedrich Hayek's attack on the feasibility of central economic planning—surely one of the least trivial contributions to one of the least trivial intellectual debates of this century. Hayek's views are summarized in his 1974 Nobel Memorial Lecture. See Friedrich August von Hayek, *The Pretence of Knowledge*, 79 AM. ECON. REV. 3 (1989) (reprint of Nobel lecture); see also F.A. HAYEK, *STUDIES IN PHILOSOPHY, POLITICS AND ECONOMICS* 27 (1967) ("The multiplicity of even the minimum of distinct elements required to produce (and therefore also of the minimum number of data required to explain) a complex phenomenon of a certain kind creates problems which dominate the disciplines concerned with such phenomena The chief difficulty . . . becomes one of in fact ascertaining all the data determining a particular manifestation of the phenomenon in question, a difficulty which is often insurmountable in practice and sometimes even an absolute one."); *id.* at 34 ("One of the chief results so far achieved by theoretical work in [the studies of society] seems to me to be the demonstration that here *individual events regularly depend on so many concrete circumstances that we shall never in fact be in a position to ascertain them all; and that in consequence not only the ideal of prediction and control must largely remain beyond our reach, . . .* The very insight which theory provides, for example, that almost any event in the course of a man's life may have some effect on almost any of his future actions, makes it impossible that we translate our theoretical knowledge into predictions of specific events. There is no justification for the dogmatic belief that such translation must be possible if a science of these subjects is to be achieved, and that workers in these sciences have merely not yet succeeded in what physics has done, namely to discover simple relations between a few observables. If the theories which we have yet achieved tell us anything, it is that no such simple regularities are to be expected." (emphasis added)); *id.* at 39 ("Perhaps it is only natural that in the exuberance generated by the successful advances of science the circumstances which limit our factual knowledge, and the consequent boundaries imposed upon the applicability of theoretical knowledge, have been rather disregarded. It is high time, however, that we take our ignorance more seriously.").

Herbert Simon's influential work on "satisficing" behavior was motivated by similar concerns. See, e.g., HERBERT A. SIMON, *THE SCIENCES OF THE ARTIFICIAL* 36 (2d ed. 1981) ("In the face of this complexity the real-world business firm turns to procedures that find good enough answers to questions whose best answers are unknowable. Thus normative microeconomics, by showing real-world optimization to be impossible, demonstrates that economic man is in fact a satisficer, a person who accepts 'good enough' alternatives, not because he prefers less to more but because he has no choice.").

range of problems, including, for example, those posed by any modern economic system,³⁷ characterized by complex interdependen-

37. Rust summarizes much recent work suggesting that a wide range of economic problems are computationally intractable. See Rust, *supra* note 32 (manuscript at 2-11).

Real economic problems are extremely high dimensional. For example if d refers to the number of goods and services in the actual economy, then d is certainly as large as one hundred thousand and probably is closer to several hundred million depending on how narrowly we define a good or service. In many problems d will also be a function of the number of agents, which is well over 5 billion if we view the planet as a single integrated economy. To the extent that we want to compute a reasonably accurate approximation to the competitive equilibrium solution . . . and for realistic values of d , the exponential complexity bounds . . . imply that an astronomical number of calculations are required. The simple logic of exponential growth tells us that . . . as d increases the number of calculations quickly grows so large that the world's fastest supercomputers would be unable to find [even] an approximate solution to the problem in any reasonable period of time.

Id. (manuscript at 10); see also David Lane & Robert Maxfield, *Foresight, Complexity, and Strategy*, in *THE ECONOMY AS AN EVOLVING COMPLEX SYSTEM II* 169, 170 (W. Brian Arthur et al. eds., 1997) (discussing the "combinatorial explosion of possible consequences" for many economic problems). Traditional economic theory provides no way out of this intractability dilemma (and does not purport to do so); as many have pointed out, conventional economic theory cures this "curse of dimensionality," and renders large-scale models analytically tractable, by assuming given and stable (rather than endogenously determined) individual preferences. See Robert H. Frank, *Positional Externalities*, in *STRATEGY AND CHOICE* 25, 44 (Richard J. Zeckhauser ed., 1991) (noting that "positional externalities"—decisions that have "important effects not only for the person who [makes them], but also for the frame of reference in which he and others operate"—are generally treated "in many economics texts as though they were an isolated exception to a normal state of affairs in which choices affect only the agents directly involved," and noting further that "the more we learn about them, the more likely it seems that actions *without* external effects may be the real exceptions"); Herbert Gintis, *Welfare Criteria with Endogenous Preferences: The Economics of Education*, 15 INT'L ECON. REV. 415, 415 (1974) (neo-classical economic theory "takes preferences as either fixed, or changing only in response to variables *external* to the model. In positive economics, the formation of preferences is relegated to sociology or social psychology; and in welfare economics, preference structures are among the fundamental, unexplained data"); John Holland, *The Global Economy as an Adaptive Process*, in *THE ECONOMY AS AN EVOLVING COMPLEX SYSTEM* 117-18 (Philip W. Anderson et al. eds., 1988) (economic choices are "determined by the interaction of many dispersed units acting in parallel [in which] [t]he action of any given unit depends upon the state and actions of a limited number of other units," requiring a "substantial extension of traditional economics"); Steffen Huck, *Trust, Treason, and Trials: An Example of How the Evolution of Preferences Can Be Driven by Legal Institutions*, 14 J.L. ECON. & ORG. 44 (1998) (reviewing notion that traditional economic theory "relies crucially on the rational actor paradigm and the assumption of given and stable preferences . . . as does the economic analysis of law"); Alan P. Kirman, *The Economy as an Interactive System*, in *THE ECONOMY AS AN EVOLVING COMPLEX SYSTEM II*, *supra*, at 491-531 (discussing the "logical and technical problems associated with standard economic models"); George J. Stigler & Gary S. Becker, *De Gustibus non Est Disputandum*, 67 AM. ECON. REV. 76 (1977) (summarizing, and justifying, economics' reliance on this assumption). This, in turn, has generated a lively debate about the realism of such models. See, e.g., JON ELSTER, *SOUR GRAPES* (1983); ROBIN HAHNEL & MICHAEL ALBERT, *QUIET REVOLUTION IN WELFARE ECONOMICS* 141-202 (1990); Jane B. Baron & Jeffrey L. Dunoff, *Against Market Rationality: Moral Critiques of Economic Analysis in Legal Theory*, 17 CARDOZO L. REV. 431 (1996); Herbert Hovenkamp, *The Limits of Preference-Based Legal Policy*, 89 NW. U. L. REV. 4 (1994); Gerald J. Postema, *Liberty in Equality's Empire*, 73 IOWA L. REV. 55, 86-87 (1987) (discussing the "so-called problem of endogenous preferences and the threat it poses to general equilibrium theory"); Amartya K. Sen, *Rational Fools: A Critique of the Behavioral Foundations of Welfare Economics*, 6 PHIL. & PUB. AFF. 317 (1977).

cies among its numerous elements. What, then, is to be done?

II. A COMPLEX SYSTEMS MODEL

To say that problems of this kind are computationally intractable is not to say that they cannot be solved; indeed, any suggestion that the Gardener's Dilemma is somehow insoluble has an odd ring to it, for gardeners seem to "solve" this problem every day. Perhaps all that we can say is that if they *are* solved—if there are efficient algorithms that enable gardeners (and others) to find progressively better configurations of these systems and perhaps even the "best" configuration, without taking an infinite amount of time to do so—we do not quite understand how that is done.³⁸ This would seem to be a matter of no small concern; if we do not understand how problems of this kind are solved, how can we choose between alternative solutions—between, say, the views of the judges of the Sixth Circuit in the majority and those in the dissent, regarding the consequences for artistic and scholarly production of re-configuring the rules of copyright?

The burgeoning study of "complex adaptive systems" begins, in a sense, with the observation that problems of this kind *are* routinely "solved" all around us. The physical and biological worlds have

38. Professor Fuller had some interesting observations about the way that "polycentric problems" of the kind we are discussing here could be solved. "So far as I can see," he wrote, "there are only two suitable methods: *managerial direction* and *contract* (or reciprocity)." Fuller, *The Forms and Limits of Adjudication*, *supra* note 30, at 398. The latter relies on "reciprocal adjustment of each center of interest with those with which it interacts," as in, for example, "an economic market [that] can solve the extremely complex problems of allocating resources, 'costing' production, and pricing goods." *Id.* at 399. As to the former, Fuller illustrated the manner in which managerial direction solves polycentric problems

by the baseball manager who assigns his players to their positions, decides when to take a pitcher out, when and whom to pinch-hit, when and how far to shift the infield and outfield for a particular batter, etc. *The relationships potentially affected by these decisions are in formal mathematical terms of great complexity—and in practical solution of them a good deal of "intuition" is indispensable.*

Id. at 398 (emphasis added). The role of human "intuition" in solving problems of this kind should certainly not be ignored; much recent work has suggested that the "parallel processing" capabilities of the human brain allow it to solve complex problems in ways that cannot be easily mimicked by ordinary computational algorithms. See, e.g., MARVIN MINSKY, *THE SOCIETY OF MIND* (1986); Rust, *supra* note 32 (manuscript at 32) (summarizing theories of cognition developed by neuroscientists and computer scientists that hypothesize that the brain's parallel processing capabilities allow it to "operate[] as a sort of 'society' or competitive economy"). But we understand little about how that intuition might operate and, regardless of the source (or even the existence) of some built-in capability to solve complex problems of this kind, intuitive solutions have the disadvantage that it is difficult, if not impossible, to choose among alternative solutions; whose intuition does one trust regarding the copyright problem discussed earlier in the text, that of the majority of judges in the Sixth Circuit or the dissenters?

innumerable examples of richly interconnected systems that, somehow, appear to reach (or more accurately, approach) optimal configurations; long-term evolutionary forces seem to have “enabled nature to ‘discover’ powerful hardware and software that is capable of solving extremely difficult large-scale computational problems that we are presently unable to solve [and that] [w]e are only just beginning to appreciate and understand.”³⁹ We believe that the study of these naturally-occurring problem-solving algorithms can shed a great deal of light on problems like the Gardener’s Dilemma and by extension, those problems in law and social policy that it closely resembles in its essential features.⁴⁰

Stuart Kauffman and his colleagues have developed a family of complex system models—known as “NK models”—for studying various problem-solving algorithms defined over complex interconnected systems.⁴¹ The model first specifies *N* (the number of system elements) and a “state space,” i.e., the number of different states (*S*) that each element can be in.⁴² Interconnections between

39. Rust, *supra* note 32 (manuscript at 3). The literature on complexity theory is vast. The best of the many readable introductions to the theory remains KAUFFMAN, *AT HOME IN THE UNIVERSE*, *supra* note 15. Other excellent introductions to the theory include PER BAK, *HOW NATURE WORKS* (1996); COMPLEXITY: METAPHORS, MODELS, AND REALITY (George A. Cowan et al. eds., 1994); PETER COVENEY & ROGER HIGHFIELD, *FRONTIERS OF COMPLEXITY: THE SEARCH FOR ORDER IN A CHAOTIC WORLD* (1995); JOSHUA M. EPSTEIN & ROBERT AXTELL, *GROWING ARTIFICIAL SOCIETIES: SOCIAL SCIENCE FROM THE BOTTOM UP* (1996); JOHN H. HOLLAND, *HIDDEN ORDER: HOW ADAPTATION BUILDS COMPLEXITY* (1995). A more technical, but readable, general treatment can be found in KAUFFMAN, *ORIGINS OF ORDER*, *supra* note 15.

40. We are not the first to suggest that the law is itself a “complex adaptive system” amenable to study through the lens of this developing body of theory. See, e.g., Thomas Earl Geu, *The Tao of Jurisprudence: Chaos, Brain Science, Synchronicity, and the Law*, 61 TENN. L. REV. 933 (1994); Andrew W. Hayes, *An Introduction to Chaos and Law*, 60 UMKC L. REV. 751 (1992); Glenn Harlan Reynolds, *Is Democracy Like Sex?*, 48 VAND. L. REV. 1635 (1995); Mark J. Roe, *Chaos and Evolution in Law and Economics*, 109 HARV. L. REV. 641 (1996); J.B. Ruhl & Harold J. Ruhl, Jr., *The Arrow of the Law in Modern Administrative States: Using Complexity Theory to Reveal the Diminishing Returns and Increasing Risks the Burgeoning of Law Poses to Society*, 30 U.C. DAVIS L. REV. 405 (1997); J.B. Ruhl, *Complexity Theory as a Paradigm for the Dynamical Law-and-Society System: A Wake-Up Call for Legal Reductionism and the Modern Administrative State*, 45 DUKE L.J. 849 (1996); J.B. Ruhl, *The Fitness of Law: Using Complexity Theory to Describe the Evolution of Law and Society and Its Practical Meaning for Democracy*, 49 VAND. L. REV. 1407 (1996) [hereinafter Ruhl, *Fitness of Law*]; J.B. Ruhl, *Thinking of Environmental Law as a Complex Adaptive System: How to Clean up the Environment by Making a Mess of Environmental Law*, 34 HOUS. L. REV. 933 (1997).

41. See KAUFFMAN, *AT HOME IN THE UNIVERSE*, *supra* note 15, at 149-89; KAUFFMAN, *ORIGINS OF ORDER*, *supra* note 15, at 40-69; Edward D. Weinberger, *Local Properties of Kauffman’s N-k Model: A Tunably Rugged Energy Landscape*, 44 PHYSICAL REV. A 6399 (1991); Edward D. Weinberger & Peter F. Stadler, *Why Some Fitness Landscapes Are Fractal*, 163 J. THEORETICAL BIOLOGY 255 (1993).

42. In what follows we will refer only to models with the simplest state space possible, i.e., two possible states for each element.

elements are modeled by assigning each element to a "spillover set," consisting of K elements.⁴³ A "fitness function" must be specified, representing each element's contribution to a system-wide variable that is to be maximized (or minimized)—aggregate "fitness" of the system as a whole—as a function of the state of the K elements in its spillover set; that is, the fitness function for each element specifies a fitness contribution for that element, given each of the S^K different configurations of the K elements in that element's spillover set.⁴⁴

One well-known algorithm for finding the fitness maximum in systems of this kind is a simple trial-and-error procedure known as the "adaptive walk."⁴⁵ An adaptive walk in a complex system proceeds as follows: Aggregate system fitness is calculated for the initial configuration in which the system begins, after which one randomly-selected element is "flipped" from state 0 to 1 (or vice versa). Aggregate system fitness is recalculated for this changed configuration, taking into account that the "flip" will affect the fitness contribution of all elements on whom the flipped element "spills over," i.e., all elements whose spillover sets include the flipped element. If system fitness post-"flip" is higher than pre-"flip"—i.e., if the new configuration has moved the system *up* the fitness landscape—we change the system configuration to the new configuration with the "flip" in place, and we repeat the process with this new configuration as the initial configuration. If, however, the change causes a decrease in system fitness—if the new configuration has moved the system *down* the fitness landscape—the change is rolled back, returning the flipped element to its starting configuration, and the process is repeated from the original configuration.

How well does this procedure identify configurations of higher fitness? The adaptive walk is quite efficient at finding the highest point on the fitness landscape in systems *that have no interconnections or spillovers* between elements (i.e., in systems with $K=1$, where an element's fitness contribution is a function *only* of its own state).⁴⁶ On

43. These K elements can be chosen at random, from among all other system elements, or in some other fashion (e.g., as those elements that are "nearby" when the elements are arrayed in a two-dimensional matrix as in Figure 1).

44. See KAUFFMAN, ORIGINS OF ORDER, *supra* note 15, at 40-69.

45. See KAUFFMAN, AT HOME IN THE UNIVERSE, *supra* note 15, at 166-80; KAUFFMAN, ORIGINS OF ORDER, *supra* note 15, at 39-40; John H. Gillespie, *Molecular Evolution over the Mutational Landscape*, 38 EVOLUTION 1116 (1984).

46. See KAUFFMAN, ORIGINS OF ORDER, *supra* note 15, at 40-46.

average, it will take no more than $N/2$ “flips”— $N/2$ steps of the adaptive walk—to find the optimal configuration of simple systems with $K=1$; the efficiency of the algorithm, in other words, scales *linearly* with the size of the system.⁴⁷

In systems with substantial spillover effects, however, the algorithm performs progressively less and less well. On these more rugged fitness landscapes,⁴⁸ the adaptive walk is increasingly likely to become trapped on *local fitness peaks*—places on the fitness landscape from which there are no steps leading upwards at all.⁴⁹ Adaptive walks on rugged landscapes, therefore, tend to end after a small number of steps when such a local peak is encountered and the walk can proceed no further.⁵⁰ That is, the more complex the system, the more completely interconnected the individual elements, the more likely it is that the adaptive walk will encounter a stable (though suboptimal) equilibrium from which no upward moves can be made, and the less likely it is that the adaptive walk will reach the global maximum, the “best” configuration for the system as a whole.

Kauffman and his colleagues have uncovered a variant of the adaptive walk—what Kauffman calls “patching”—that appears to represent an improvement over the adaptive walk algorithm in these more complex systems.⁵¹ The patching algorithm operates by assigning each element in an NK system to a single group of elements, or “patch.”⁵² See Figure 2. The adaptive walk begins, as before, by “flipping” the state of a randomly selected element. The algorithm

47. See *id.* at 46.

48. See *id.* at 47.

49. At the limit, i.e., in systems with maximally rugged fitness landscapes caused by random spillovers from all elements to all other elements, the number of such local optima increases *exponentially* with N . See *id.* (“these extremely rugged landscapes are so rife with local optima that trapping on such optima is essential [sic] inevitable”); see also Stuart Kauffman & Simon Levin, *Towards a General Theory of Adaptive Walks on Rugged Landscapes*, 128 J. THEORETICAL BIOLOGY 11 (1987) (describing relationship between inter-element spillover, landscape “ruggedness,” and the effectiveness of adaptive walk algorithms).

50. On randomly constructed systems with spillovers among all elements, Kauffman has shown that adaptive walks terminate, on average, after the $\ln(N)$ th step. See KAUFFMAN, *ORIGINS OF ORDER*, *supra* note 15, at 48-49. Thus, in such a system consisting of 100 elements, although the state space is enormously large (consisting of 2^{100} configurations), average walk length is only approximately $\ln(100)=4.61$ steps long. See *id.* In addition, the relative “height” of these local peaks gets lower and lower, on average, as the interconnections among elements increase and cause fitness landscapes to become more rugged, a phenomenon Kauffman labels the “complexity catastrophe.” See *id.* at 52-54; KAUFFMAN, *AT HOME IN THE UNIVERSE*, *supra* note 15, at 194-95.

51. See KAUFFMAN, *AT HOME IN THE UNIVERSE*, *supra* note 15, at 252-67; see also Ruhl, *Fitness of Law*, *supra* note 40, at 1469-72 (describing Kauffman’s patching algorithm and speculating on its implications for the study of legal systems).

52. See Ruhl, *Fitness of Law*, *supra* note 40, at 1469-72.

then calculates the effect of this flip on the aggregate fitness of the members of the patch of which the flipped element is a member, rather than, as before, determining whether the new system configuration increases fitness for the system as a whole.⁵³ If patch fitness post-“flip” is higher than pre-“flip,” we change the system configuration to the new configuration with the “flip” in place (regardless of whether the “flip” causes the fitness of the system as a whole to increase or decrease) before continuing the adaptive walk; if patch fitness post-“flip” is lower than before, the change is rolled back (again, regardless of the effect of the change on fitness of the system as a whole), returning the flipped element to its starting configuration, before moving on.⁵⁴

FIGURE 2
The system depicted in Figure 1, divided (“geographically”) into four patches of 25 elements each.

What the patching algorithm does, in other words, is to evaluate any change of state *solely* in terms of the effects of the change on the fitness contribution of members of an element’s patch. Patching is the equivalent of the following decision rule: Individual elements will be permitted to move from one state to another if, but only if, the effect of the move on the aggregate fitness of the members of its patch is a positive one. If members of the element’s spillover set are not

53. See *id.*
54. See *id.*

also members of the element's patch, the move from one state to another may affect the fitness of many individuals outside of its patch. Those effects, however, will be entirely ignored when deciding whether or not the move is to be permitted; that will be evaluated solely on the basis of the *within-patch* effects of the change.⁵⁵ The patching algorithm seeks local, within-patch improvements in fitness rather than global improvements; instead of adopting state changes that impact positively on the system as a whole, it adopts state changes that impact positively on some subset of the system, some decision-making entity that can accept or reject the proposed move depending solely on the effects of that move on the entity's aggregate fitness. Each patch is allowed to maximize its own fitness, independent of any effects on the fitness of non-members or on the aggregate fitness of the system as a whole.⁵⁶

Kauffman has shown that this patching algorithm can, in certain circumstances, dramatically increase the efficiency of the search for high aggregate system fitness; an adaptive walk over a patched system finds, in a given number of steps, higher points on the fitness landscape *for the system as a whole* than the same walk over the same system without patching.⁵⁷ At least in certain circumstances, allowing sub-system decision-making units to accept or reject state changes of its members can be an effective means for moving the entire system across the fitness landscape to find peaks of high global fitness. Letting each of these subsets seek its own optimum can be an efficient way to find the optimum for the system as a whole.

The result is a fairly remarkable one:

It is by no means obvious that the [highest aggregate fitness] of the [system] will be achieved if [it] is broken into quilt patches, each of

55. *See id.*

56. Understanding the distinction between "patches" and "spillovers" is essential for understanding the patching algorithm. Patches and spillovers each represent connections between elements, though they function in very different ways. An element's "spillover set" consists of those elements whose fitness contribution is a function of that element's state; anthropomorphizing a bit, the members of my spillover set are those individuals whose "fitness contributions" change as a result of my action. An element's "patch," on the other hand, consists of those elements whose fitness contributions must increase if the element is to be permitted to take any particular action, the elements that have some *voice* in determining whether or not other patch members are allowed to change their state. An individual's "spillover set" and "patch" may consist of the same elements, the two sets may have some partial overlap, or the two sets may be entirely disjoint; as we discuss below, a measure of how disjoint the two sets are—what we call "congruence"—turns out to be an important determinant of the efficiency of this algorithm for finding high points on the fitness landscape. *See infra* pp. 1081-84.

57. *See* KAUFFMAN, *AT HOME IN THE UNIVERSE*, *supra* note 15, at 256-64.

which tries to [maximize its own fitness] regardless of the effects on surrounding patches. Yet this is true. *It can be a very good idea, if a problem is complex and full of conflicting constraints, to break it into patches, and let each patch try to optimize, such that all patches coevolve with one another.*⁵⁸

Patching is effective, it appears, because it reduces the tendency of the adaptive walk to become trapped on suboptimal local fitness peaks. As noted above, the simple adaptive walk often becomes trapped on such local peaks, unable to dislodge itself (which requires moving *down* the landscape) in order to reach higher points beyond. Patching ameliorates this problem by allowing such downward moves; the patching algorithm allows the system to adopt configurations that yield lower aggregate fitness (provided that the change increases local patch fitness), which serves to dislodge the system from otherwise stable but suboptimal fitness peaks. By moving to points lower on the fitness landscape, it can reach other, higher configurations that it would otherwise be unable to reach under the simple adaptive walk. It is, in other words, precisely the systemically *destabilizing* effects of the patching procedure that makes it effective.

Kauffman's description of patching as quite possibly a "fundamental mechanism underlying adaptive evolution in ecosystems, economic systems, and cultural systems"⁵⁹ is an apt one and generates a number of important questions about the way that complex problems can be solved in a wide range of systems. For example, what are the characteristics of systems in which patching is or is not effective? What makes for a "well-designed" patch for problem-solving purposes? If patching works by destabilizing the search for optimal configurations of complex systems, dislodging the search from suboptimal local fitness peaks, can this effect be *too* destabilizing, preventing steady hill climbing over the fitness landscape? Are there general rules pursuant to which patches can be formed that render them more or less effective for the task of solving

58. *Id.* at 262 (emphasis added); *see also id.* at 264 ("[H]ard problems with many linked variables and loads of conflicting constraints can be well solved by breaking the entire problem into nonoverlapping domains."). Professor Ruhl describes the patching algorithm as follows:

[T]ake a hard, conflict-laden task in which many parts interact, and divide it into a quilt of nonoverlapping patches. Try to optimize within each patch. As this occurs, the couplings between parts in two patches across patch boundaries will mean that finding a "good" solution in one patch will change the problem to be solved by the parts in adjacent patches. [C]hanges in each patch will alter the problems confronted by neighboring patches, and the adaptive moves by those patches in turn will alter the problem faced by yet other patches

Ruhl, *Fitness of Law*, *supra* note 40, at 1469.

59. KAUFFMAN, AT HOME IN THE UNIVERSE, *supra* note 15, at 264.

for fitness maxima?

We have begun building NK systems to explore these questions, with particular attention to the way that the construction of patch boundaries affects the effectiveness of the patching algorithm. Consider again a patched NK system such as the one displayed in Figure 2. Each element in such a system can be said to be a member of two different groups: a "spillover group" (consisting of those elements whose contribution to aggregate fitness is affected by the state of the element in question) and a "patch" (consisting of those elements whose aggregate fitness is measured to evaluate whether the element's change of state is to be allowed). These two sets may be entirely disjoint or completely overlapping (or something in between these two extremes); that is, the relationship between patch membership and spillover in any system may be such that the effects of an element's change of state are felt *only* by members of *other* patches, by a mixture of elements some of whom are and some of whom are not members of its patch, or entirely by members of its own patch.

We can thus define, for any NK system, a variable—what we call "congruence"—to measure the overlap between patches and spillover sets. Congruence is defined as the proportion of elements in an individual's spillover set that are also members of that individual's patch, averaged over all elements in the system.⁶⁰ Congruence can

60. For example, imagine a system of 10 elements ($N=10$), with four elements in each element's spillover set ($K=4$), divided into two patches of 5 elements each. Patch 1 contains elements 1, 2, 3, 4, and 5, Patch 2 the other elements (6-10). Spillovers are distributed as follows (an X in the row i column j in the following table indicates that element i 's spillover set includes element j):

	-----Patch 1 elements-----					-----Patch 2 elements-----					
	1	2	3	4	5	6	7	8	9	10	Congruence
1	X	X	X	X							1.00
2		X	X	X	X						1.00
3			X	X	X	X					0.75
4				X	X	X	X				0.50
5					X	X	X	X			0.25
6						X	X	X	X		1.00
7							X	X	X	X	1.00
8	X							X	X	X	0.75
9	X	X							X	X	0.50
10	X	X	X							X	0.25

Element 3's spillover set thus consists of elements 3, 4, 5, and 6, and its congruence is thus 0.75 (because 3 of the 4 members of its spillover set are also members of Patch 1). The average congruence for this system is 0.70. Notice that, in this example, the fitness contribution of element 6 will be affected by a change in element 3's state, but that effect will play no role in the

vary from zero (representing a system in which patch and spillover sets are entirely disjoint for all elements and, thus, a system in which any element's change of state affects *only* the fitness contribution of members of *other* patches) to one (representing a system in which patch and spillover sets overlap completely and, thus, a system in which all spillover effects are "internalized" within patches where the change of any element's state affects *only* members of its own patch).

What happens to the effectiveness of the patching algorithm as this measure of within-patch spillover internalization moves from zero to one? We have constructed several thousand NK systems, systematically varying the congruence of each to observe the effect of those variations on the ability of the patching algorithm to find points of high aggregate system fitness. Details of our modeling procedures are set forth in the Appendix; our results are summarized in the graphs displayed in Figure 3.⁶¹

decision whether to keep or to "roll back" a change in element 3's state inasmuch as element 6 is not a member of element 3's patch (Patch 1).

61. The complete data set is available from the authors.

Figure 3

A view of representative results of our adaptive walk simulations. Each point on the graphs represents the result of a single 10,000 step adaptive walk over a 1000-element system. Values of percentage change in aggregate system fitness are plotted on the Y-axis; system congruence is plotted on the X-axis. Results are displayed for systems defined with patches of 33, 40, 45, 50, and 66 members, and with three different values for K (the extent of inter-element spillover) (K=8, K=12, and K=24).

Each point in the graphs in Figure 3 represents the results of a single adaptive walk of 10,000 steps over a 1000-element system; the graphs show the percentage increase in aggregate system fitness at the conclusion of the adaptive walk on the Y-axis and system congruence on the X-axis. Figure 3 displays results for systems constructed with relatively high inter-element spillover ($K=24$), intermediate spillover ($K=12$), and low spillover ($K=8$) at each of five different patch sizes (33, 40, 45, 50, and 66).

As the graphs strikingly demonstrate, the effectiveness of patching as a search algorithm—the ability of the adaptive walk to uncover higher points on the system's fitness landscape—is a positive function of system congruence; generally speaking, it is apparent that more congruent systems are able to uncover higher points on the fitness landscape than less congruent systems. At all values of K and at all patch sizes, the curves appear to have roughly similar shapes, showing increasing search efficiency with increasing degrees of congruence, although in all cases, interestingly, search efficiency appears to level off or drop at the *highest* levels of system congruence.⁶²

III. IMPLICATIONS: PATCHING, CONGRUENCE, AND THE THEORY OF COMPETITIVE FEDERALISM

To summarize, we have suggested in the foregoing that (a) the problem of finding "better" configurations of complex interconnected systems is both a highly general one and one that is in many circumstances analytically intractable, (b) allowing sub-groups ("patches") to seek their own optimal *local* configurations can improve the efficiency of searching for *global* optima, and (c) the effectiveness of patching as a means of solving these complex

62. These qualitative results are confirmed by a more systematic quantitative analysis. We analyzed our entire data set, comprising 11,686 separate adaptive walks, by means of multiple linear regression, using percentage system fitness change (our proxy for search efficiency) as the dependent variable, and K , patch size, and congruence as the independent variables. The results indicate that almost 90% of the variation in search efficiency can be explained by variation in these three independent variables ($R^2=0.877$); the best-fit least-squares equation for these data is

$$Y (\% \text{ change in fitness}) = 0.07 - 0.002 (K) + 0.00004 (\text{Patch size}) + 0.27 (\text{Congruence})$$

Only the relationship with congruence and K are statistically significant; that is, holding K and patch size constant, variations in congruence explain a highly statistically significant amount of the variation in search efficiency ($t=265.73$, $P<<0.001$), as do variations in K (holding patch size and congruence constant) ($t=-34.79$, $P<<0.001$). Patch size, however, is more weakly (though still statistically significantly) related to variations in search efficiency ($t=2.19$, $P=0.025$).

problems is dependent upon the relationship between patch boundaries and spillover effects between individual elements—i.e., patching appears to work best in systems with the “right” balance between those inter-element effects that are internalized *within* patches and those that are not.

Reflecting the focus of this symposium, we believe that our analysis has important implications for questions of Internet governance and legal theory. Legal institutions are (or should be) designed to solve problems defined over complex systems—problems like the simple “coursepack” problem discussed above. If we are to have effective problem solving in this complex policy space, a central goal for the design of legal institutions is the formation of congruent, independently optimizing decision-making sub-groups.

This is not a new idea. Processes whereby “complex systems” are divided into “patches,” with each “patch” allowed to permit or prohibit “changes of state” of “patch members” on the basis of the aggregate “within-patch” effects of that change on all patch members, are familiar enough in political and legal realms. We know them as decentralized, federalist rule-making systems, systems in which individuals are members of non-overlapping groups, each of which decides whether or not to permit or prohibit particular conduct based upon the effects of that conduct on the perceived welfare of the individuals within the sub-group, and they are, of course, ubiquitous.

They also tend to define their “patch boundaries” geographically. At the global level, for instance, the boundaries around legally-recognized decision-making patches—sovereign states in the international system—have, at least since the Treaty of Westphalia in 1648, been defined in geographical terms; the power and the right to enact and enforce law are lodged with sovereigns whose domains of authority are circumscribed largely in *territorial* terms, applicable within a particular geographic area to persons (real or fictional), events, and conduct located there.⁶³

This is, of course, no accident. Both the ability to ascertain the will of a proximate people and the ability to control the actions of

63. See references cited in Post, *The “Unsettled Paradox”*, *supra* note 1; see also RESTATEMENT (THIRD) OF FOREIGN RELATIONS LAW § 201 (1987) (“Under international law, a state is an entity that has a defined territory and a permanent population, under the control of its own government, and that engages in, or has the capacity to engage in, formal relations with other such entities.”). This pattern, of course, is repeated at lower levels of the governance hierarchy, e.g., among states or provinces within national federal systems, municipalities or counties within states, etc.

those people by means of enforceable rules stem from traditional reliance on communications technologies that operate best at close range and on enforcement by means of physical forces that work best over shorter distances. And the foregoing discussion suggests other ways in which this "makes sense." Defining boundaries of decision-making patches in geographic terms is likely to be an effective mechanism to find the highest "fitness peak" for any system *where spillovers are themselves distributed geographically*. That is, given a system in which spillover effects are physically clustered and attenuate with increasing distance from the locus of the participating individuals, a system in which "neighbors" (in the *geographical* sense) are more likely to affect the welfare of each other than they are to affect the welfare of those less physically proximate, dividing such a system into geographically-defined patches will produce largely congruent decision-making units, allowing more effective searching for the system-wide optima.

We think it reasonable to suggest that inter-individual spillover effects in human societies have, for the bulk of human history, been distributed largely in this manner, that the effects of individual conduct on the welfare of others have been determined to a significant extent by the geographical proximity between them, and that the impact of individual actions has been felt (as a general matter) most intensely by those closest, *physically*, to the actor and/or to the transaction. Defining patch boundaries—the boundaries between sovereign states and the units within those states—in geographical terms makes for effective searching in policy space in such a world.

But the assumption of geographically-clustered spillovers no longer holds nearly as powerfully as it once did; indeed, the course of human history can be described in terms of a slow, but now rapidly-accelerating increase in the magnitude of "between-patch spillovers"—an increase in the likelihood that the effects of conduct (state changes) in one geographically-defined jurisdiction (patch) will have substantial effects on large numbers of individuals from other patches. Systems like these, in which an orderly relationship between patch membership and spillovers has been perturbed in a substantial way, must find ways to "re-couple" these two parameters—to re-establish the congruence of decision-making groups—if they are to continue to function as efficient problem-solving mechanisms. As the Internet helps to complete the transition from largely geographical to

largely a-geographical spillover effects—as spillover effects between individuals become distributed more-or-less at random with respect to the geographical location of those individuals—*continuing to define patch membership in geographical terms produces decision-making groups of low congruence, and the independent decisions of such patches will be increasingly unlikely to find high peaks on the global welfare landscape*.⁶⁴

This is the central problem of Internet governance: how to constitute the decision-making units that can, through the “pull and tug” of contending rule-sets, continue to find configurations of rule-sets that move the global interconnected system “upwards” in welfare-space. How should the boundaries of these units be set? This would be a difficult enough task even were this perturbation the result of a single systemic “shock” that had somehow altered the pattern of spillover effects; but the Internet poses an even more difficult problem, for the spillover patterns are themselves shifting and doing so in “Internet time” as a consequence of the rapid evolution of the underlying communications technologies themselves.⁶⁵ What process or processes are capable of adjusting

64. Readers should note that we are *not* suggesting that cyberspace is somehow a “self-enclosed regime” where communications have no consequences for individuals or entities not participating in those communications. See Jack L. Goldsmith, *Against Cyberanarchy*, 65 U. CHI. L. REV. 1199, 1215-16 (1998) (suggesting that “regulation skeptics” like us somehow fail to recognize that “[a] communication in cyberspace often has consequences for persons outside the computer network in which the communication took place” and that “communications via cyberspace produce harmful, real-world effects on those who have not consented to the private ordering of the cyberspace community”). Our argument, rather, is that the Internet and other related communications technologies cause those spillover effects from cyberspace to real space to be distributed more or less at random with respect to physical geography.

65. See Bernard J. Hibbitts, *Yesterday Once More: Skeptics, Scribes and the Demise of Law Reviews*, 30 AKRON L. REV. 267, 314 n.182 (1996) (“The shift from institutions to individuals as the primary locus of publishing activity may only be facilitated by the speed at which Internet publishing technology evolves. Already journalists are talking about ‘Internet time’. [sic] the accelerated rate at which new Web browsers and other applications are developed and enter the marketplace. Change is occurring so fast that it is almost impossible for actual or would-be publishing organizations to keep up: by the time they collectively decide to take one step, the technology has advanced by another. In this situation, autonomous individuals may be the only agents able to make consistent use of the latest publishing innovations.”); Moisés Naím, Editor’s Note, 107 FOREIGN POL’Y 5, 5 (1997) (pace of change “overwhelms the capacities of most individuals and institutions to grasp all of its implications or realize all of its interconnections”); Bert Ely, *Financial Services Modernization*, REGULATION, Spring 1997, at 18, 19, *available online* (last visited Nov. 19, 1998) <<http://www.cato.org/pubs/regulation/reg20n2-per.html>> (arguing that modernization efforts in banking regulation “are proceeding against the backdrop of technological changes that make the current regulatory regime unworkable. Technology is making it increasingly difficult to neatly compartmentalize banking, insurance, and securities products.”); Michele Matassa Flores, *Free Range Celebrates Elder Status on the Web*, SEATTLE TIMES, Apr. 18, 1996, at E1 (“Everything about the Internet is changing so fast, there should be a way to interpret Internet ‘time’ in the same way we translate a mutt’s age into dog years.”);

patch boundaries to "track" these rapidly changing spillover patterns?

At the very least, would-be regulators of cyberspace need to pay close additional attention to the ways that these two parameters can be re-coupled. We have suggested elsewhere that the Internet calls for a higher degree of deference to rulemaking within a-geographical, decentralized, voluntary associations,⁶⁶ and we believe that the foregoing provides normative underpinnings for this view. Allowing individuals to define the boundaries of their own, a-geographical patches by voluntary movement into and out of decision-making bodies that have little or even no ties to a particular physical location⁶⁷—what we might call "self governance"—may allow both more rapid and more "congruent" responses to shifts in spillover patterns. Individuals are more likely to be in possession of the relevant information regarding the effect of spillover on their own welfare and can act more quickly on that information than can agents at a higher level of the organizational hierarchy (i.e., their elected representatives) to whom that information must be re-directed and

John Markoff, *A Quicker Pace Means No Peace in the Valley*, N.Y. TIMES, June 3, 1996, at D1 (quoting Andrew Grove, chief executive of the Intel Corporation, "we are now living on Internet time" as "traditional product cycles give way to an endless stream of upgrades"); Kathy Rebello, *Inside Microsoft*, BUS. WK., July 15, 1996, at 56 ("Microsoft, already the ultimate hardcore company, is entering a new dimension. It's called Internet time: a pace so frenetic it's like living dog years—each jammed with the events of seven normal ones."); see also *Bensusan Rest. Corp. v. King*, 126 F.3d 25, 27 (2d Cir. 1997) (analogizing application of established law to the "fast-developing world of the [I]nternet" to "trying to board a moving bus").

66. See Johnson & Post, *Law and Borders*, *supra* note 1, at 1387-1402 (arguing that responsible a-geographical self-governance institutions can emerge on the Internet); Johnson & Post, *The New "Civic Virtue"*, *supra* note 1, at 25-26 ("The best available solution to conflicts in individual goals and values regarding online conduct may be found by allowing individuals to join distinct, bounded communities on the Internet, each with its own divergent set of rules, and by allowing those communities to deal with *external* pressures by devising their own mechanisms for filtering out unwelcome messages and with *internal* conflict by easing (or requiring) exit. Democratic debate and traditional legislative action may not, after all, be the best way to make the best public policy for the Internet. If we can preserve individual liberty to make educated and empowering choices among alternative online rule sets, our most thoughtful and high-minded collective-action option may be to abandon the process of elections and deliberations regarding some single best law to be imposed impartially on all from the top down. We may instead find a new form of civic virtue by allowing the governance of online actions to emerge 'from the bottom up' as a result of the pull and tug between local online 'jurisdictions' that do not attempt to act in a dispassionate or disinterested or 'public-spirited' manner.").

67. The notion of a-geographical rule-making entities is itself not a novel one; we are all embedded in a web of decision-making bodies—fictional legal persons (e.g., partnerships or corporations), voluntary associations of all kinds (e.g., the Catholic Church, the shareholders of General Motors, or the members of Usenet discussion groups)—that exist on a plane separate from, and often entirely independent of, geographical sovereigns. All may have their own decision-making processes for determining whether particular conduct on the part of its members is, or is not, permitted. The boundaries around these groups can be geographically fluid, determined by the interaction between each association's "membership" rules and the individuals' decisions to move into or out of the group.

re-processed before "official" boundary realignment can occur. Individuals on whom spillover is most directly concentrated can, using their ability to "enter" and "exit" such venues, have the most impact on the rules applicable to the spaces from which the spillover emanates, which may allow both more rapid and more "congruent" responses to shifts in spillover patterns. In other words, putting boundary-definition in the hands of individuals directly affected by the rules made within those boundaries may allow a faster and more flexible response to rapidly changing spillover patterns.⁶⁸

68. We do not mean to suggest that this exhausts the possible universe of "recoupling" devices; indeed, we would hope that the analysis presented here would stimulate work on constructing a typology of such devices. As an example of the kinds of devices we have in mind, consider suggestions that have been advanced for solving spillover/patch "uncoupling" in the context of metropolitan area governance. As in cyberspace, a central problem for metropolitan area governance in the world of atoms is precisely the increasing degree to which spillover effects and the boundaries between municipalities have become uncoupled, *see, e.g.*, Richard Briffault, *The Local Government Boundary Problem in Metropolitan Areas*, 48 STAN. L. REV. 1115, 1133 (1996) (observing that local boundaries "probably always generated some spillovers, but in the past, when local governments were set farther apart by unincorporated land, and people focused more of their activities within the territorial limits of their particular locality, the spillovers may have been relatively slight The spillover problem is more acute today because local borders frequently abut one another, and people range widely in their daily activities across multiple local boundaries. In contemporary metropolitan areas, local governments are sure to generate externalities and area residents are sure to be excluded from participating in the decisions of many localities that have direct implications for their lives."), and Professors Ford and Frug have proposed a novel recoupling mechanism in this context. Given existing patterns of extensive cross-border spillover, they argue, the boundaries of local patches (municipalities) must shift to take those spillovers into account; we have to "replace our current legal conception of localities with one that embraces the ageographical city" by "treat[ing] people not as located solely in one jurisdiction but as 'switching center[s]' for all the networks of influence' within the region[s] that affect their lives." Jerry Frug, *Decentering Decentralization*, 60 U. CHI. L. REV. 253, 323 (1993) (second alteration in original) (footnote omitted). This can be accomplished, they suggest, by means of changes in local voting rules to allow nonresidents to vote in local elections. Local elections would be

"open to all members of a metropolitan region or even to all citizens of a state." All local elections in the region would be held on the same day; voters would receive a number of votes equal to the number of local offices to be filled and could cast them wherever they chose. Under such a system of regionwide cumulative voting for local office, "voters would effectively draw their own jurisdictional boundaries, decide which local governments were most important to them, and allocate their votes accordingly." . . . With political rights decoupled from residency, the need to deny localities the power to engage in locally self-interested or exclusionary regulation disappears.

Briffault, *supra*, at 1156-57 (summarizing Ford and Frug proposals) (second emphasis added) (quoting Richard Thompson Ford, *The Boundaries of Race: Political Geography in Legal Analysis*, 107 HARV. L. REV. 1843, 1909-10 (1994) [hereinafter Ford, *The Boundaries of Race*]); *see also* Richard Thompson Ford, *Beyond Borders: A Partial Response to Richard Briffault*, 48 STAN. L. REV. 1173 (1996) [hereinafter Ford, *Beyond Borders*]; Ford, *The Boundaries of Race*, *supra*, at 1857-59 (proposing regionwide cumulative voting); Frug, *supra*, at 329-30 (same).

The premises underlying this solution to the boundary problem are, we believe, largely correct: that individuals may be in the best position to process and act upon information about the way that actions taken within particular decision-making venues affect their own personal utility and that they can be relied upon to act accordingly, i.e., to join (and thereby influence the

As for the implications of this work for legal theory more generally, it might come as little surprise to those familiar with theories of "competitive federalism" that dividing up a complex system into independent self-optimizing decision-making patches can increase the efficiency of the search for optimal system-wide configurations.⁶⁹ Those theories reflect a broad consensus regarding two underlying benefits of decentralized decision-making procedures of this kind.⁷⁰ First, decentralized decision-making can function as an efficient sorting mechanism; mobile individuals can, via migration into (and exit out of) different decision-making units, efficiently match their preferences for different menus of local public goods.⁷¹ Second, dividing a decision-making polity into smaller, local decision-making sub-units may, because information about local conditions and local preferences is imperfectly distributed and tends to be concentrated locally, be subject to fewer inefficiencies of information transfer; therefore, local governments and consumers will be more likely to make better (welfare-maximizing) decisions.⁷² There is an equally broad consensus on the "cost" side of the decentralization

rules within) the relevant decision-making patch. See Ford, *Beyond Borders*, *supra*, at 1189 n.45; Frug, *supra*, at 329-30.

69. Good introductions to the vast literature on competitive federalism include Vicki Been, "Exit" as a Constraint on Land Use Exactions: Rethinking the Unconstitutional Conditions Doctrine, 91 COLUM. L. REV. 473 (1991); Dan L. Burk, *The Market for Digital Piracy*, in BORDERS IN CYBERSPACE, *supra* note 4, at 205; John D. Donahue, *Tiebout or Not Tiebout? The Market Metaphor and America's Devolution Debate*, J. ECON. PERSP., Fall 1997, at 73, 73-82; Frank H. Easterbrook, *Antitrust and the Economics of Federalism*, 26 J.L. & ECON. 23 (1983); Richard A. Epstein, *Exit Rights Under Federalism*, 55 LAW & CONTEMP. PROBS. 147 (1992); Robert P. Inman & Daniel L. Rubinfeld, *Rethinking Federalism*, 11 J. ECON. PERSP. 43 (1997); Wallace E. Oates & Robert M. Schwab, *Economic Competition Among Jurisdictions: Efficiency Enhancing or Distortion Inducing?*, 35 J. PUB. ECON. 333 (1988); Susan Rose-Ackerman, *Does Federalism Matter? Political Choice in a Federal Republic*, 89 J. POL. ECON. 152 (1981); Barry Weingast, *The Economic Role of Political Institutions: Market-Preserving Federalism and Economic Development*, 11 J.L. ECON. & ORG. 1 (1995).

70. See generally Yingyi Qian & Barry R. Weingast, *Federalism as a Commitment to Preserving Market Incentives*, 11 J. ECON. PERSP. 83, 83 (1997) (describing these "two well-known sources of benefits from decentralization").

71. This derives from the well-known Tiebout model. See Charles M. Tiebout, *A Pure Theory of Local Expenditures*, 64 J. POL. ECON. 416 (1956). Tiebout demonstrated that an idealized system comprising (a) a "perfectly elastic" supply of jurisdictions, (b) perfectly mobile inhabitants, (c) perfect information about the characteristics of each jurisdiction, and (d) no inter-jurisdictional externalities can produce the optimal level of public goods within the system as a whole. See Dan L. Burk, *Virtual Exit in the Global Information Economy*, 73 CHI.-KENT L. REV. 943 (1998) and references cited therein.

72. This theme derives from the work of Friedrich Hayek. See HAYEK, *supra* note 36, at 66-81 (suggesting that a "polycentric order" without a "directing center" can be superior to a "hierarchical order" by dispensing with necessity of communicating all the information on which its several elements act to a common center); F.A. Hayek, *The Use of Knowledge in Society*, 35 AM. ECON. REV. 519 (1945).

equation as well: decentralized decision-making is disfavored where jurisdictions are not "congruent," i.e., where there are significant intercommunity interdependencies or spillovers.⁷³

A greater understanding of the patching algorithm described here—greater understanding, we cheerfully admit, than the authors of this paper now have—may shed some light on these mechanisms. Patching may be more than merely a metaphor for decentralized political decision-making structures (though it is that and no less interesting because of it); those structures may, in a sense, be instantiations of the patching algorithm in the political realm. Federalism may "work," in other words, because it is a "patching algorithm," a means for solving public policy problems defined over a most complex "social welfare landscape."⁷⁴ As such, an understanding of the factors that determine the effectiveness of the algorithm cannot help but have an impact on our understanding of these political decision-making institutions.

73. It is virtually axiomatic, in the law and economics literature, that intercommunity spillovers distort the otherwise efficient outcomes reached by decentralized decision-making institutions. See, e.g., ROBERTA ROMANO, *THE GENIUS OF AMERICAN CORPORATE LAW* 5 (1993) (noting that while the "benefits from federalism are axiomatic in American politics, it is also well recognized that a federal system can . . . diminish individual welfare . . . if the costs and benefits of a specific public policy do not fall within the boundaries of one jurisdiction"); Frank H. Easterbrook & Daniel R. Fischel, *Mandatory Disclosure and the Protection of Investors*, 70 V.A. L. REV. 669, 697 (1984) (observing general rule, in connection with interstate competition for corporate charters, that inter-jurisdictional competition is most effective "when the consequences of a decision will be experienced in one jurisdiction"); Epstein, *supra* note 69, at 149 (noting that the inter-jurisdictional spillover problem is not addressed by a federalist system relying on exit rights to maintain efficient allocations of public goods); Clayton P. Gillette, *In Partial Praise of Dillon's Rule, or, Can Public Choice Theory Justify Local Government Law?*, 67 CHI.-KENT L. REV. 959, 971 (1991) (discussing general notion that "[o]utside the Tiebout world of no externalities . . . some constraints on local power are necessary to prevent strategic local behavior"); Inman & Rubinfeld, *supra* note 69, at 46-47 (observing that "when there are significant intercommunity interdependencies (like pure public goods or spillovers)," dispersed decision-making by decentralized sub-units "may result in economically inefficient public policies," and noting that "[f]or most economists," dispersed decision-making with "a strong central government to provide pure public goods and control intercommunity externalities, essentially defines what federalism is about"); Robert P. Inman & Daniel L. Rubinfeld, *The Political Economy of Federalism*, Center for the Study of Law and Society, University of California Berkeley School of Law (Working Paper No. 94-15, 1994) (discussing the assumption of the absence of inter-jurisdictional externalities in the Tiebout model and the central importance of understanding the extent of spillover effects in determining the appropriate decision-making units); cf. Richard L. Revesz, *Rehabilitating Interstate Competition: Rethinking the "Race-to-the-Bottom" Rationale for Federal Environmental Regulation*, 67 N.Y.U. L. REV. 1210, 1250-52 (1992) (observing that "race-to-the-bottom" arguments in regard to inter-jurisdictional competition are distinct from arguments that interstate externalities may justify federal intervention).

74. See Ruhl, *Fitness of Law*, *supra* note 40, at 1469-73 (interpreting the common law and the federalist division of decision-making authority in the U.S. Constitution as patching devices).

Our analysis, for example, may help illuminate the role of inter-jurisdictional spillover in patched systems. There is a suggestion in our data that the efficiency of dispersed decision-making processes is not a simple inverse function of the magnitude of inter-jurisdictional spillovers, i.e., that configuring the boundaries between jurisdictions in such a way that all inter-jurisdictional externalities are "internalized" is not a necessary condition for the effective functioning of the patching algorithm. While it is true that systems with high congruence appear to be more efficient at finding system-wide fitness peaks than those with more inter-patch spillover, there appears to be somewhat more to the story. In the systems we have thus far examined, perfectly congruent systems with no inter-group externalities are often *less* effective at finding system-wide optima than systems with a somewhat lower degree of congruence. See Figure 3. In other words, our results suggest—tentatively, to be sure, but provocatively—that search efficiency may decline if congruence is too low *or* too high and that there may be an intermediate congruence "sweet spot" causing systems with an intermediate level of congruence to be better at finding higher points on the fitness landscape than both those in which spillovers are only weakly internalized within patches (low congruence) *and* those in which spillovers were *perfectly* internalized within patches (high congruence).

This interpretation is consistent with our understanding of the underlying mechanics of the patching algorithm. As noted above, patching appears to be effective precisely because it is *destabilizing*. It allows local configurations to change in ways that may be sub-optimal in the short term from the standpoint of the system as a whole, driving the system *down* from suboptimal foothills in fitness space; but these moves alter the environment of other local units, generating reactions and adjustments by these adversely affected "neighbors" and creating a pull and tug among conflicting rule sets that ultimately allows the overall matrix to achieve a better solution over the course of a large number of moves. To anthropomorphize, at the group level, patching facilitates a kind of conversation among conflicting patches, a dialogue that may improve the search for better configurations overall. Justice Brandeis memorably praised federalism as a means to allow "a single courageous state [to] serve as a laboratory . . . without risk to the rest of the country";⁷⁵ it may well

75. *New State Ice Co. v. Liebmann*, 285 U.S. 262, 311 (1932) (Brandeis, J., dissenting).

be, however, that it is of some systemic value that some "local experiments" *do* pose risks to other jurisdictions, causing those jurisdictions to confront (and to solve) new problems that permit new frontiers of the fitness landscape to be explored.

CONCLUSIONS

[L]ike many other precious human goals, the rule of law may best be achieved by not aiming at it directly. What is perhaps most needed is not an immediate expansion of international law, but an expansion of international community, multiplying and strengthening the bonds of reciprocity among nations. When this has occurred—or rather *as this occurs*— . . . the law can act as a gardener who prunes an imperfectly growing tree in order to help the tree realize its own capacity for perfection. This can occur only when all concerned genuinely want the tree to grow and to grow properly. Our task is to make them want this.⁷⁶

We are confident that the study of behavior in the world of bits will, perhaps paradoxically, illuminate the role that physical and geographical constraints play in the world of atoms because only by eliminating geography can we truly understand the role that geography plays in human affairs.⁷⁷ We have admittedly traveled somewhat far afield for the underpinnings of the preceding discussion, but we believe that Fuller was (on this and many other questions) largely correct—the law's role is to "prune[] an imperfectly growing tree in order to help the tree realize its own capacity for perfection."⁷⁸ We need to take the Gardener's Dilemma and the limitations of our ability to comprehend the movement of complex systems across rugged landscapes more seriously. We should, perhaps, concentrate less on trying to specify optimal configurations of legal systems and more on the design of processes through which more, rather than less, favorable configurations may emerge, less on the question "what is the best rule to govern a particular transaction?" and more on the question "what is the best algorithm for finding more acceptable rule configurations?" At the very least, an understanding of generally applicable principles that describe the behavior of interdependent systems of all kinds should inform our

76. Fuller, *Adjudication and the Rule of Law*, *supra* note 30, at 8.

77. This is, perhaps, not paradoxical at all. As the legal anthropologist Leopold Pospisal was fond of pointing out, if we lived in a world that was entirely of one color—red, say—we would not only be unable to truly understand 'green' or 'blue,' we would be unable to understand 'red' itself.

78. Fuller, *Adjudication and the Rule of Law*, *supra* note 30, at 8.

policy debate in cyberspace and elsewhere.

The proof, of course, is in the pudding; the value of this (or any other) conceptual model lies in its ability (or lack thereof) to provide us with insight into the behavior of actual systems and to generate interesting second- and third-order hypotheses about observed phenomena. We fully recognize that our work has only scratched a very (dare we say it?) complex surface and that further investigation is needed if the full implications of this approach to social systems are to be uncovered. We hope that our efforts stimulate others to help us in the search for a more thorough and refined understanding of the behavior of these systems.

APPENDIX: SIMULATING ADAPTIVE WALKS IN COMPLEX SYSTEMS

The computer-based NK systems that we studied were built using a program written by Jeff Williams in Java (the source code for which is available from the authors). The program has an outer "shell," through which the user can set up five parameters defining the NK system, and an inner "engine" that generates the system and runs the adaptive walk algorithm. The five parameters are:

- The number of elements (N);
- The number of elements in each element's spillover set (K);
- The number of equal-sized patches into which elements will be placed (P);
- The number of "steps" of the adaptive walk (i.e., the number of "flips" of randomly chosen elements before the search terminates); and
- The desired value for system congruence (C , $0 \leq C \leq 1$).

Once the user chooses values for each of these parameters,⁷⁹ the program generates an NK system by first constructing a "pool" of N elements and assigning each of the N elements to one of the P patches by subdividing the N -element pool into groups of the appropriate size, depending upon the chosen values for N and P ; that is, to generate P equally-sized patches in an N -element system, each patch must contain N/P members, and the program therefore assigns the first N/P elements to Patch 1, the next N/P elements to Patch 2, and so on for all P patches.

The program chooses the K elements in each element's spillover set at random from among the other elements of the system, subject to two important constraints. First, each element is assumed to be a member of its *own* spillover set; that is, we have embodied the plausible (but not logically necessary) assumption that the state of every element always affects its *own* fitness. Inasmuch as we were primarily interested in studying the way that congruence influenced the effectiveness of patching and the adaptive walk, we designed the program to choose the remaining $K-1$ elements so as to achieve the pre-set level of congruence (C) for *each* element (and, therefore, for the system as a whole). In order to achieve a congruence of C , an

79. The program allows the user to specify a range for each of these parameters, allowing the program shell to select random values from the ranges for each run.

element's spillover set must contain $(C \cdot K)$ of its patch members. Because each element is, by definition, a member of its own patch, and because each element has already been placed into its own spillover set, an additional $(C \cdot K) - 1$ *within-patch* elements must be included in that element's spillover set. Each patch contains N/P elements and, therefore, there are $(N/P) - 1$ other elements (besides the individual element whose spillover set is being constructed) in the patch; the program chooses $(C \cdot K) - 1$ of these patch members at random and places them into the spillover set. Having thus chosen $(C \cdot K)$ members of the spillover set, the program chooses the remaining elements $(K - (C \cdot K))$ at random from among the other non-patch elements.

By way of illustration, consider a system with the following parameters:

- $N = 1000$
- $K = 8$
- $P = 20$
- $C = 0.75$

Constructing the spillover set for any element—element number 51, say—proceeds as follows. The system is first divided into 20 patches (of 50 elements each); elements 51-100 are placed into patch 2. To achieve a congruence of 0.75, 6 of the 8 elements of element 51's spillover set must be members of patch 2. Element 51 is placed into its own spillover set first; 5 other elements are chosen randomly from patch 2 and added to element 51's spillover set. The remaining two members of element 51's spillover set are chosen at random from all system elements *other* than those in patch 2.

Completing this process for each of the N elements produces an $N \times N$ *spillover table* showing the elements in each element's spillover set; the spillover table for our hypothetical system might look as follows, with an "X" in row i column j indicating that element j has been placed into element i 's spillover set:

	1	2	3	4	...	51	...	998	999	1000
1	X									
2		X		X		X				
3			X					X		
4				X						X
...										
51						X			X	
...										
998			X					X		
999	X								X	
1000						X				X

This table shows, for example, that elements 51 and 999 are in element 51’s spillover set. The full 1000x1000 spillover table would show K “Xs” in each row, representing the K elements in each element’s spillover set.

From this spillover table the program constructs the converse “spill-in” table for each element. That is, just as row 51 shows all elements on whom element 51 spills over, column 51 shows all elements that “spill in” to element 51, all elements whose state helps determine the fitness contribution for element 51. (In the above example, we see that elements 2, 51, and 1000 spill over onto element 51 and are therefore members of element 51’s “spill-in” set.)

Element 51’s fitness contribution in any system configuration is computed as follows. The fitness contribution of element 51 is, by

definition, a function of the states of the K elements in its spill-in set.⁸⁰ There are 2^K possible configurations for element 51's spill-in set (since each element in the spill-in set can be in one of two possible states). In the absence of any other information regarding the nature of the interconnections between the elements, the fitness contribution for element 51 given any particular configuration is assigned at random;⁸¹ that is, the fitness table for element number 51 is constructed by assigning a randomly chosen fitness value to each possible state configuration of the elements in the spill-in set.⁸²

The adaptive walk begins in a randomly chosen system configuration. The program first computes the fitness contribution of each individual element. This requires (a) examining the state (1/0) of each of the members of that element's spill-in set and (b) consulting the fitness table to determine the element's fitness contribution. The program stores both the aggregate fitness for each patch (the sum of the fitness contributions of each patch element) and for the system as a whole.

A randomly chosen element is then "flipped" from its current state to the opposite state (i.e., from 0 to 1 or vice versa). The program re-computes the fitness contribution of all elements whose contribution may be affected by the flip, i.e., all elements in the spillover set of the flipped element. The program then re-computes patch fitness for the flipped element's patch (i.e., the sum of the fitness contributions of all elements belonging to that patch). If that sum is greater than (or equal to) the fitness of that patch prior to the flip, the adaptive walk continues from the changed configuration; if

80. On average, given that we have assigned elements to spillover sets at random, each element's spill-in set will similarly contain K elements.

81. Kauffman describes the rationale for random assignment as follows:

Several factors suggest that we might fruitfully model complex [inter-element] interactions by assigning the effects *at random* in some way. First, we may as well admit our ignorance in the biological cases at the present moment. Second, we are trying to build rather general models . . . to begin to understand what such landscapes look like and what organismal features bear on landscape ruggedness. If we model the fitness effects of [inter-element] coupling "at random," we will obtain the kind of general models of landscapes we seek. Third, if we are lucky, we will find that real landscapes in some cases look very much like our model landscapes. . . . In short, we can seek general laws in biology concerning the structure and even the evolution of fitness landscapes by building models.

KAUFFMAN, AT HOME IN THE UNIVERSE, *supra* note 15, at 171.

82. In the illustration above, given 8 elements in element 51's spill-in set, the fitness table for element 51 will contain 256 (2^8) entries, one for each of the 256 combinations of these 8 elements. To avoid unnecessary storage, in actual design this table starts empty and is constructed as needed. If a requested entry does not exist, a new random value is created and stored in the table.

that sum is less than patch fitness prior to the flip, the flip is "rolled back" (i.e., the flipped element is returned to the state it was in prior to the flip), and the walk continues.

The process terminates when the program has completed the specified number of steps. The program's output consists of the difference between the aggregate system fitness at the conclusion of the adaptive walk and aggregate system fitness in the initial configuration.

All NK systems analyzed in this paper contained 1000 elements ($N=1000$) and the adaptive walks were allowed to proceed for 10,000 steps ($S=10,000$). The program's outer shell varied the other parameters as follows:

- K (size of spillover set) was randomly chosen from the set $\{4, 6, 8, 10, 12, 24\}$.
- P (number of patches) was randomly chosen from the set $\{12, 13, \dots, 33\}$.
- C (the desired value for system congruence) was randomly chosen from the set $\{1/K, 2/K, \dots, K/K\}$.