

University of San Diego

## Digital USD

---

University of San Diego Law and Economics  
Research Paper Series

Law Faculty Scholarship

---

May 2005

## The Web of Law

Thomas A. Smith

*University of San Diego Law School*, [tacsmith@sandiego.edu](mailto:tacsmith@sandiego.edu)

Follow this and additional works at: [https://digital.sandiego.edu/lwps\\_econ](https://digital.sandiego.edu/lwps_econ)



Part of the [Law and Economics Commons](#)

---

### Digital USD Citation

Smith, Thomas A., "The Web of Law" (2005). *University of San Diego Law and Economics Research Paper Series*. 8.

[https://digital.sandiego.edu/lwps\\_econ/art8](https://digital.sandiego.edu/lwps_econ/art8)

This Article is brought to you for free and open access by the Law Faculty Scholarship at Digital USD. It has been accepted for inclusion in University of San Diego Law and Economics Research Paper Series by an authorized administrator of Digital USD. For more information, please contact [digital@sandiego.edu](mailto:digital@sandiego.edu).

**THE WEB OF LAW****Thomas A. Smith<sup>↑</sup>**

The usefulness of the World Wide Web is largely due to its network structure. Nearly a billion web pages - the number grows daily - are connected through “hyperlinks” that allow one to go quickly from one web document to another. In spite of its enormous size, the Web is a “small world.” One can “surf” from any web page to any other in only about twenty clicks, on average. This small world property, which the Web shares with many other natural and artificial networks, also accounts for the “six degrees of separation” in the world’s social network. Surprisingly enough, every person on the planet is linked to every other through, on average, only about six friends or acquaintances.<sup>1</sup>

In part because of intriguing properties like these, scientists and mathematicians in recent years have become intensely interested in the structure of networks. Networks turn out to be crucial to understanding everything from physics and biology,<sup>2</sup> to economics<sup>3</sup> and sociology. This article proposes that the science of networks has important contributions to make to the study of law as well. While network science is beginning to be applied in the physical, biological, and social sciences, legal scholars have yet to study, or even recognize as

---

<sup>↑</sup> © 2004 Thomas A. Smith. Professor of Law, University of San Diego. Very special thanks to the LexisNexis Corporation, without whose very generous help and skills, this article would not have been possible. Very special thanks to Larry Solum for his many substantive contributions and especially for his unflagging encouragement to pursue this article. Thanks to Bob Hillman for his suggestions and tough questions, which led me to count citations by hand, and then realize node aging must be occurring. I acknowledge a debt of \$5 to Bob. Thanks to Patrick and Luke Smith, who helped me count hundreds of citations by hand, even if I did pay them for it. This article is dedicated to TI and to the memory of my late colleague Paul Wolmouth, who I think would have liked it.

<sup>1</sup> See generally Duncan J. Watts, *SMALL WORLDS: THE DYNAMICS OF NETWORKS BETWEEN ORDER AND RANDOMNESS* 1999; DUNCAN J. WATTS, *SIX DEGREES: THE SCIENCE OF A CONNECTED AGE* 2003.

<sup>2</sup> Albert-László Barabási\* & Zoltán N. Oltvai, *Network Biology: Understanding the Cell’s Functional Organization*, *NATURE REVIEWS GENETICS* 5, 101-113 (2004).

<sup>3</sup> Souma, Wataru; Fujiwara, Yoshi; Aoyama, Hideaki *Complex networks and economics*, *PHYSICA A*, Volume 324, Issue 1-2, p. 396-401.

such, one of the largest, most accessible, and best documented human-created networks in existence. This is the centuries-old network of case law and other legal authorities into which lawyers, judges, and legal scholars routinely delve in order to discover what the law is on any given topic.<sup>4</sup>

To see that American<sup>5</sup> case law is a network, observe how it resembles the Web in structure. The system of American case law, however, is not just *like* a network. It *is* a network. It has the peculiar mathematical and statistical properties that networks have. It can be studied using techniques that are now being used to describe many other networks, some found in nature, and others created by human action.<sup>6</sup> Studying the legal network can shed light on how the legal system evolves, and many other questions. To initiate what I hope will become a fruitful new type of legal scholarship, I present in this article the preliminary results of a rudimentary but significant citation study of nearly four million American legal precedents, which was undertaken at my request by the LexisNexis corporation using their well-known Shepard's citation service. This study demonstrates that the American case law network has the overall structure that network theory predicts it would.


To return to the Web/legal system analogy, recall that Web pages are linked to each other through "hyperlinks," on which the web surfer clicks to be taken automatically to the web page referred to by the hyperlink. Cases, similarly, are linked to one another by citations. Just as one can explore the Web by moving from one web page to another through hyperlinks, so one can move through case law by following the trail of citations. The term "web" derives, of course, from the web-like structure created by web pages and hyperlinks. One can create a picture of the Web, representing each web page as a point—a "node" or "vertex" in network science terminology—and each connecting hyperlink as a line—a "link" or "edge" in network

---

<sup>4</sup> I plan to avoid in this article murky jurisprudential questions about what the relationship is between the interconnected set of legal authorities (cases, statutes, and so forth) and the thing we call "the law." However, I will note that while this relationship may amount to less than an identity, there is still some sense in which all these authorities and the relationships among them documented by citations are a very significant *part* of the law.

<sup>5</sup> Not just American, of course. I begin with American only for the sake of convenience. Legal networks extend from jurisdiction to jurisdiction, but the links between countries are fewer between than within countries and other jurisdictional entities. The significance of the density of links is discussed *infra* at \_\_\_\_\_. The analysis in this article could probably be applied to any common law system. To the extent networks are generated by decisions in civil law systems, it could be applied to them as well.

<sup>6</sup> ADAM FERGESON, AN ESSAY ON THE HISTORY OF CIVIL SOCIETY 1767, Pt. 3.2 ("... nations stumble upon establishments, which are indeed the result of human action, but not the execution of any human design."). Hayek developed this insight with his notion of "spontaneous order." *See generally* BOUDEWIJN BOUCKAERT, ANNETTE GODART-VAN DER KROON, HAYEK REVISITED \_\_\_\_\_ 2000.



**Figure 1: Graph of a portion of the World Wide Web.** Points from which links emanate are webpages (nodes), and the lines between such points are hyperlinks (links or edges in graph theory parlance). Courtesy of [Scientific American].

Terminology. *Figure 1* above is a partial graph of the Web. Web pages are points, and hyperlinks are lines branching out from those nodes, which connect them to other nodes. A graph of the network of case law would look similar to *Figure 1* in many respects. The network structure of legal authority is not just a curiosity. Different types of networks have different properties. Their topology can be studied rigorously. These properties affect how the network organizes itself over time. The realization that American case law forms a network, a mathematical object of a particular kind, suggests hypotheses about the overall and internal structure of the system of legal authorities that can be tested empirically.

This article has three parts. First, I introduce some basic concepts of network science, including such important ideas as nodes, links, random graphs, evolving networks, scale-free networks, small worlds, the “rich get richer” dynamic, node fitness, and clusters. Oddly enough, the mathematical tools that have proven most useful for studying networks (or at least scale-free networks) come from statistical mechanics, a branch of physics.<sup>7</sup> While much of the mathematics in play is unfortunately beyond the skills of most legal scholars, including this author, most of the important ideas are not. Readers already familiar with basic network theory

---

<sup>7</sup>Reka Albert & Albert-Laszlo Barabasi, *Statistical mechanics of complex networks*, 74 REVIEWS OF MODERN PHYSICS, Jan. 2002.

may wish to skip to Part II, but others will find the theory described in Part I both accessible and, I hope, more than mildly interesting.

In Part II, I present evidence that demonstrates that there is a high likelihood that the network of American case law is a “scale-free network,” a particularly interesting kind of network described by physicist Albert-Laszlo Barabasi,<sup>8</sup> Reka Albert, and others,<sup>9</sup> who have studied the structure of the Web. This evidence is the result of a study generously conducted by LexisNexis at my request.<sup>10</sup> Using a program written for the purpose, LexisNexis determined the citation frequency of all of the nearly four million federal and state cases in its database, and the citation frequency of all U.S Supreme Court cases.<sup>11</sup> They then graphed the citation frequencies to determine whether, as I hypothesized, they would conform to a power-law distribution, as they would if American case law constituted a scale-free network. The power-law distribution of a scale-free network was conspicuously evident, indicating that, in all likelihood, American case law constitutes a scale-free network.<sup>12</sup> Indeed, the resulting graphs can fairly said to be “classical,” so closely do they resemble in form similar graphs produced for other networks, such as the web.<sup>13</sup> To those familiar with the network science literature, this result in a citation study will not be surprising. However, to legal scholars who have never thought of the system of legal authorities as a “scale-free network,” the empirical demonstration that it is one should be significant in its own right. The study is also novel in that it is, while rudimentary, the largest citation study, in number of citations, to my knowledge, ever performed.<sup>14</sup>

---

<sup>8</sup> R. Albert, H. Jeong, & A.-L. Barabási, *Diameter of the World Wide Web* NATURE 401, 130-131 (1999); Albert-László Barabási & Réka Albert, *Emergence of scaling in random networks* SCIENCE 286, 509-512 (1999).

<sup>9</sup> The study of scale-free networks in particular and complex networks in general has generated a large literature. A helpful bibliography of network-related literature generally may be found at the website of the Santa Fe Institute. See *Santa Fe Institute, Complex Interactive Networks, Bibliography* [http://discuss.santafe.edu/dynamics/stories/storyReader\\$8](http://discuss.santafe.edu/dynamics/stories/storyReader$8)

<sup>10</sup> gratitude

<sup>11</sup> [They have also provided me with the data and software that will enable me to measure citation frequency of the sets of cases decided by virtually any American state or federal court.]

<sup>12</sup> Other possible explanations?

<sup>13</sup> See *infra* at \_\_\_\_\_.

<sup>14</sup> S. N. DOROGOVTSSEV, J. F. F. MENDES, *EVOLUTION OF NETWORKS: FROM BIOLOGICAL NETS TO THE INTERNET AND WWW* \_\_\_\_\_ 2003 contains a summary of studies of networks for degree distribution, including citation networks. The nearly four million node U.S. legal citation network analyzed in the Smith/LexisNexis study appears to be the largest by number of nodes by some margin. However, this study measures only “in degree,” a fairly rudimentary measure. Part of the purpose of this paper is to suggest to network scientists that the legal network is well worth their attention.

Having introduced network theory in Part I, and having presented evidence that American case law is a scale-free network in Part II, I argue for the significance of this discovery in Part III. I hope that by the time they reach Part III, readers will already be realizing the potential richness of applying network theory to legal systems. In Part III, I describe some insights that appear from this application and suggest areas for future research.

The most famous hypothesis about the structure of law is that it is a “seamless web.”<sup>15</sup> This old phrase, however, is just a metaphor we have used to grope for a reality we have not been in a position to express more precisely. Network science changes that. Law is a web, and now we can describe precisely features of its structure we could only guess at before. For example, if law is a scale-free network, it is not a seamless web, but a fragmented one. It will have organized itself spontaneously into clusters of cases tightly connected among themselves, but only loosely connected to other clusters. These clusters will be semantically significant. How law has organized itself may, or may not, correspond to the traditional legal categories lawyers and scholars are familiar with. By studying the network structure of law, we can discover what the natural organization of our legal system is—natural in the sense that it has not been imposed on law by scholars or officials, but has developed over time by virtue of the millions of connections judges<sup>16</sup> have made among cases in deciding them. Other legal systems may also be studied from a network perspective. They may turn out to have topologies similar to, or different from, that of our own. Either result would shed light on whether legal systems share architectural similarities and whether and where they differ.

Scale-free networks have hub and spoke structures, like those of the airline route maps at the back of glossy in-flight magazines. A few hub cities, such as Chicago, New York, and Atlanta, get most of the links. This is true of the legal network as well. Links (citations) are distributed approximately according to a power law.<sup>17</sup> This means that by far most cases are

---

<sup>15</sup> Apparently, the origin of this phrase is unclear. Ethan Katsch explains in \_\_\_\_\_ that

There is considerable ambiguity about the origin of this expression. Frederic Maitland, an English legal historian, appears to have been the first to use the phrase “seamless web” in a law-related context. Maitland wrote: “Such is the unity of all history that any one who endeavours to tell a piece of it must feel that his first sentence tears a seamless web.”

Ethan Kahn, \_\_\_\_\_ fn 3, *citing* Frederic William Maitland, *A Prologue to a History of English Law*, 14 L. QUARTERLY REV. 13 (1898); *see also* 1 FREDERICK POLLOCK & FREDERIC W. MAITLAND, *THE HISTORY OF ENGLISH LAW* 1 (2d ed. 1899).

<sup>16</sup> Or judges, scholars and legislators, to the extent they cite other legal authorities.

<sup>17</sup> Power-law distributions of case law are shown *infra* at \_\_\_\_\_.

rarely cited and, once decided, disappear into oblivion. Relatively few cases are cited many times. What determines whether a case makes it into the elite of cases that are cited hundreds or thousands of times, instead of just a few, or never?<sup>18</sup> By studying the statistical dynamics of citation over time, scholars using network theory could shed significant light on what accounts for the success of a legal authority. How much, for example, does a case's merely being decided earlier account for its citation frequency? How much of a difference does a case having been decided by a higher court make? What about the "fitness" of a case, in terms of its persuasiveness or analytical acuity? Do these attributes explain in part a case's flourishing, survival, or extinction as a precedent? Do cases have a natural life span? Does their authority tend to wax and wane, and does this depend on the type of case?<sup>19</sup> It is possible that if we study the evolution of the legal network, the "Web of Law," we will discover unsuspected historical dynamics in legal authority. Perhaps the characteristics of legal evolution themselves have changed over time. If there have been changes in the dynamics of legal evolution, or other noticeable changes, perhaps they correspond to recognized watersheds in legal history. As I discuss briefly below, using network theory to analyze law may enable us to understand, in a much more rigorous way than previously possible, the dynamics by which interpretations of important laws, such as the Constitution and landmark statutes, change over time. By analyzing the legal network, we can shed light on these questions and many more. Network science offers important new opportunities for the empirical study of legal systems.

## I. INTRODUCTION TO NETWORK THEORY

A network is just a set of items, which we term *nodes* or *vertices*, with connections between them, termed *links* or *edges*. Networks are mathematical objects, but there are concrete examples everywhere. There are social networks of friends and acquaintances, economic networks of producers and customers, and scientific networks of research collaborators. Networks also abound in nature. Blood vessels form a distributional network,

---

<sup>18</sup> Legal scholars will also want to ask the same question about law review articles. Their citation frequency probably follows a power-law distribution as well. This study would be relatively easy to conduct with the program LexisNexis has developed. Legal scholars will also want to ask the same question about law review articles. Their citation frequency probably follows a power-law distribution as well.

<sup>19</sup> See *infra* at \_\_\_\_\_. I suspect that some aspects of the shape of the citation frequency curve of the Web of Law may be explained by the "aging" of authorities.

the brain contains a neural network, and predators and prey form food networks in ecosystems, to name just a few examples.<sup>20</sup>

### A. From Random Graphs to Scale-Free Networks

Mathematicians have studied networks, in the form of graph theory, at least since \_\_\_\_\_ Euler's solution of the Königsberg Bridge problem in 1735.<sup>21</sup> The renaissance of contemporary graph or network theory can be dated to 1959, when Erdos and Renyi began publishing a series of eight important papers on random graphs.<sup>22</sup> Random graphs are different from the networks we are concerned with, but they are a good place to start.<sup>23</sup> To construct a simple random graph, let us begin with fifteen nodes. We number the nodes, 1 through 15. Then we select two nodes at random, say 6 and 8, and establish a link between them. We carry on this procedure for some specified number of links, each time picking randomly the two nodes to be connected.<sup>24</sup>




Figure 2. (a). A Erdos-Renyi random graph.

(b). A Poisson distribution of link degree.

Erdos and Renyi proved that in a large random graph, each node will have approximately the same number of links, or, to use network theory terminology, each node will

<sup>20</sup> R. J. Williams, E. L. Berlow, J. A. Dunne, A.-L. Barabási, and N. D. Martinez, *Two degrees of separation in complex food webs*, PROCEEDINGS OF THE NAT'L ACADEMY OF SCIENCES **99**, 12913-12916 (2002).

<sup>21</sup> Euler, L., 1736. *Solutio problematis ad geometriam situs pertinentis*. Commentarii Academiae Scientiarum Imperialis. Petropolitanae 6, 128-140 (in Latin). A nice, brief summary of the history of network theory may be found in [L.A.N. Amaral, J.M. Ottino, *Complex systems and networks: challenges and opportunities for chemical and biological engineers*, CHEMICAL ENGINEERING SCIENCE **59** (2004) 1653 - 1666.]

<sup>22</sup> Erdős, P. & Rényi, A. *Publ. Math. Inst. Hung. Acad. Sci.* **5**, 17-61(1960).

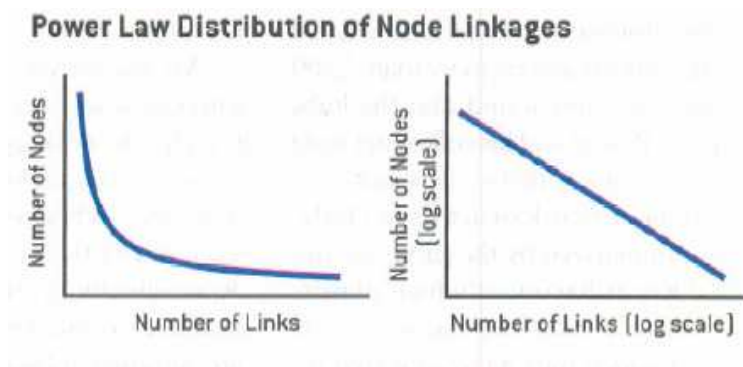
<sup>23</sup> Random graphs not simple.

<sup>24</sup> [http://www.nd.edu/~networks/PDF/NatureGen\\_Vol5-Feb04\\_Barabasi.pdf](http://www.nd.edu/~networks/PDF/NatureGen_Vol5-Feb04_Barabasi.pdf)



be of approximately the same “degree.”<sup>25</sup> Link degree in a random graph follows a Poisson distribution,<sup>26</sup> which resembles the familiar bell curve. The mean of this distribution is obviously a meaningful statistic. It tells us a lot about any particular random network, such as how many links a typical node has, and how “interconnected” the network is.

Many important networks, however, are not random. Instead of nodes having approximately the same number of links, a few nodes have many links, while most nodes have only a few. This structure emerged dramatically when Albert-Laszlo Barabasi and \_\_\_\_\_ Albert studied the network structure of the Web. At the beginning of their study, they expected to find that the number of links web pages have followed a Poisson distribution, as in a random graph.<sup>27</sup> Instead, they found that the number of links followed a power-law distribution.



**Figure 3.** Scale-free networks have a power-law distributions of nodes

Barabasi and Albert coined the term “scale-free network” to describe a network of this kind, alluding to the fact that the mean number of links did not meaningfully describe the scale of the network as the means of many other distributions do.<sup>28</sup> In a scale-free network, the node with the mean degree is not typical. Barabasi and his colleagues set out to discover why the Web had a degree distribution so strikingly different from the Poisson distribution familiar

<sup>25</sup> ALBERT-LASZLO BARABASI, LINKED: THE NEW SCIENCE OF NETWORKS 2002.

<sup>26</sup> Background on Poisson distribution

<sup>27</sup> <http://www.nd.edu/~networks/PDF/Scale-Free%20Sci%20Amer%20May03.pdf>

<sup>28</sup> A more rigorous and better explanation may be found in WATTS, supra note \_\_\_\_\_ [Six Degrees] at 107. As he explains, if one were to graph a normal distribution in a log-log format as power-law distributions normally are, then the curve would be a line with a definite  $x$  and  $y$  intercepts and concave to (bulging outward from) the origin. One could say these definite “cut offs” at the  $x$  and  $y$  axes define the “scale” of the network that has a normal distribution (which classical random graphs do not, but let us suppose some network does have a normal distribution of some attribute). A network with a power-law distribution of in-degree, however, such as scale-free networks have, does not have any definite  $x$  or  $y$  intercepts, and so could be said to be in this sense “scale free.”

from random graphs. They discovered an elegant explanation that hinges on two features that distinguish scale-free networks, such as the Web, from random networks.

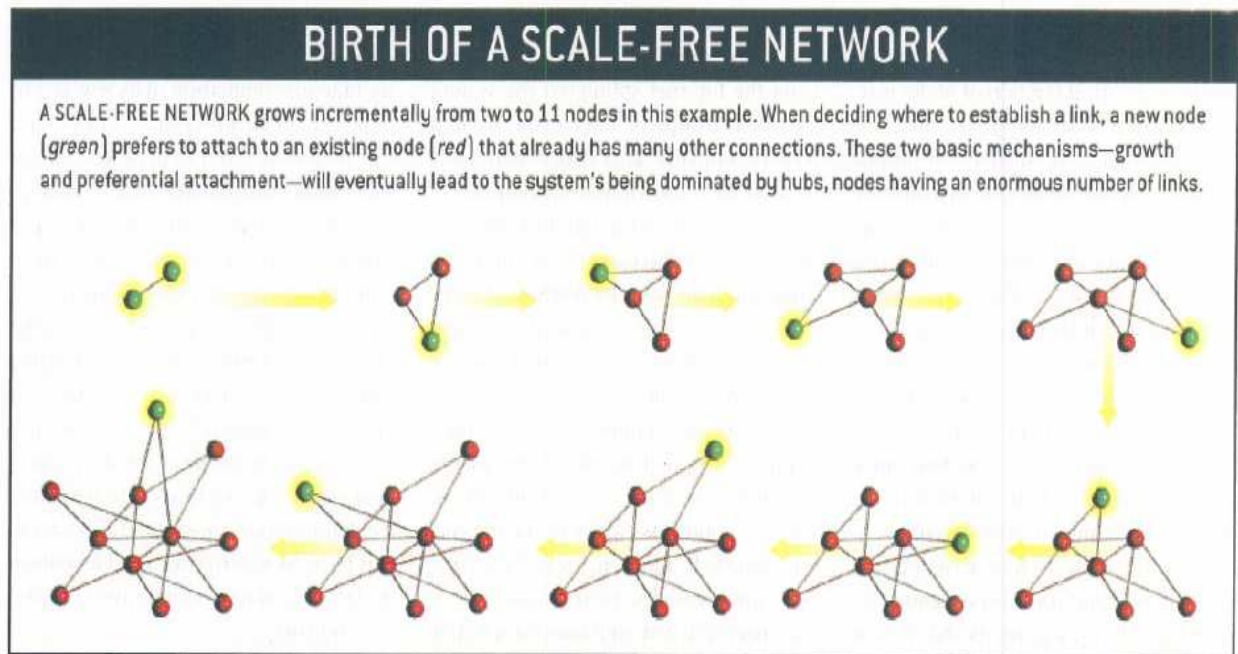
Recall that random networks begin with a fixed number of nodes. Links are then added to connect randomly selected nodes. This means that each node in the random network has a certain probability of being chosen as one of the next pair of nodes that will get a new link added between them. In the Web, however, and in many other both human-created and natural real networks, the number of nodes is not fixed. New web pages are created every day, and in many other networks as well, the number of nodes is constantly growing. This is the first important distinction between classical random networks and scale-free networks. Scale-free networks are created by a dynamic process, in which the number of nodes grows over time.<sup>29</sup>

The second important feature of scale-free networks is that new links are not added randomly. In classical random networks, each node has the same probability of getting a new link added to it. The probability that a node in a scale-free network will acquire a new link, however, depends on how many links it already has.<sup>30</sup> In modeling scale-free networks, Barabasi and Albert made the probability of acquiring new links a function of how many links a node already had. Nodes with more links had a greater chance than nodes with fewer links of acquiring additional links as the network grew. In terms of links, the “rich got richer.” Barabasi and Albert called this mechanism “preferential attachment.” Barabasi and his colleagues were able to prove that the combination of these two features, network growth and preferential attachment, produced the scale-free network structure, with its power-law distribution of node degree.

---

<sup>29</sup> This is not the only way a scale-free network can come to exist. It is, however, an important class of scale-free networks, and includes, in all likelihood, the Web of Law.

<sup>30</sup> In this article, I focus on the so-called Barabasi-Albert (“BA”) model of the generation of scale-free networks. However, since Barabasi and Albert published their influential paper, network scientists have described other processes that give rise to scale-free networks. Nevertheless, the BA model is probably the most prominent mechanism described in the literature for generating a scale-free network, and, more importantly, it provides an attractive combination of plausibility and simplicity as a model for the generation of the legal network. Thus in this article, I concentrate the BA model. In Part III.\_\_\_\_ below, however, I address the quite plausible possibility that the legal network is an *aging* network, like the network of scientific citations. Casual inspection of the curves of legal citation frequency compared to other networks, suggests that it well may be. Aging networks depart from the scale-free structure in predictable ways, and in ways that casual inspection suggests may account for the shape of the citation frequency curve of the Web of Law. If so, considering the Web of Law as an aging network suggests many interesting consequences for our understanding of precedent and the doctrine of stare decisis in our legal system. *See infra* at \_\_\_\_\_.



**Figure 4.** Scale-free networks form when networks grow over time and links are formed by preferential attachment. From Barabasi, *Scientific American*.

This was a significant discovery in network science. Prior to Barabasi and Albert's work on the structure of the Web, the study of network mathematics<sup>31</sup> was dominated by the Erdos-Renyi random graph model. However, many, if not most, of the networks of interest to us are not random. People do not pick new friends or sexual partners<sup>32</sup> at random, economic actors do not pick new suppliers, customers, or employees at random, and judges do not cite cases at random. Barabasi and other scientists had discovered a fundamental self-organizing process within real world networks.

### ***B. Small World Networks***

A striking feature of scale-free networks is their "small world" quality.<sup>33</sup> Regarding the Web, which has almost a billion web pages, for example, one might think that to get from one particular web page to another would take many clicks. If a billion nodes were arranged in a circle, and one chose two nodes at random, it would be very improbable that those two nodes would be within a million nodes of each other. Stepping from one node to the next in the circle, it would probably take millions of steps to get from one randomly chosen node to another. Yet, one can get from any web page to any other on the Web in, on average, about twenty clicks. The "diameter" of the Web is thus much smaller than one might guess. Sociologists discovered the same phenomenon by studying social networks.<sup>34</sup> The mere "six degrees of separation" said to lie between every person in the world, is perhaps the original example of the small world phenomenon.<sup>35</sup>

Short cuts make large, spread-out networks "smaller," that is, they reduce the average distance between nodes, as *Figure 4* below illustrates. Thus a Chinese peasant might seem thousands of steps away or more from a U.S. Senator, in terms of somebody who knows somebody.<sup>36</sup> Yet, social short-cuts may greatly reduce the number of intermediaries the peasant would have to go through to contact the Senator. Suppose a Chinese foreign exchange student knows both a Chinese rural provincial governor and a U.S. Senatorial intern. These links may

---

<sup>31</sup> But a network can be a small world without being scale free. Cf. sociology

<sup>32</sup> Fredrik Liljero, Christofer R. Edling, Luis A. Nunes Amaral, Sexual networks: implications for the transmission of sexually transmitted infections, *MICROBES AND INFECTION* 5 (2003) 189-196.

<sup>33</sup> Duncan J. Watts, *Small World*, Princeton U. Press 1999; M. Kochen (Ed.), *The Small World*, Ablex, Norwood, N.J., 1989.

<sup>34</sup> Watts, *Six Degrees*

<sup>35</sup> Watts, *Six Degrees*

<sup>36</sup> And they might be: this is average separation only.

greatly reduce the social distance. Not just social networks, but many different sorts of networks have this small world quality.<sup>37</sup> Indeed, small world networks are ubiquitous.

Small world networks have many interesting properties. For example, the “diameter” of a network, that is, the average minimum number of links in the shortest path from any node to any other, provides an indication of how interconnected and integrated a network is. The diameter of a network might also change as it grows. This can tell us whether a network is getting more or less integrated over time, and whether there are limits as to how integrated it will become.

The small world phenomenon was originally discovered and described by sociologists,<sup>38</sup> and social networks provide some of the most intuitive examples of small worlds.





FIG. 15. The random rewiring procedure of the Watts-Strogatz model, which interpolates between a regular ring lattice and a random network without altering the number of nodes or edges. We start with  $N=20$  nodes, each connected to its four nearest neighbors. For  $p=0$  the original ring is unchanged; as  $p$  increases the network becomes increasingly disordered until for  $p=1$  all edges are rewired randomly. After Watts and Strogatz, 1998.

**Figure 5.** From Albert and Barabasi, *Statistical Mechanics of Complex Networks*.

<sup>37</sup> <http://www.nytimes.com/library/national/science/061698sci-smallworld.html>

<sup>38</sup> Milgrom etc.




*Figure 6.* A high school social network, showing clusters. From James Moody, *Race, school integration, and friendship segregation in America*, 107 *AMERICAN JOURNAL OF SOCIOLOGY* 679-716 (2001).

Consider, for example, a social network that is organized, as most are, into cliques or clusters that, to use Watts's term, are weakly connected to other cliques. In *Figure 6* above, high school friendships are depicted as a network. The yellow nodes are white students, the green, black students, and the red, students of other races. Lines represent friendship links. Note that students with no friends, or only one friend, are outside or on the periphery of the network. Those with many friends are more centrally located. While there are numerous links connecting students of different races, it is apparent that individuals are still clustered by race. Each racial cluster has two sub-clusters [defined by gender]. The links that cross through the central area of the graph are some of the short cuts that make this social network a small world. Because of the short cuts, a white male student might have to go through only one or two other students to meet a female black student, or vice versa. Without these intermediaries, the student would have to go through many more people around the perimeter of the network. Shortcuts such as these make the high school social network a much "smaller world" than it would otherwise be. How "small," or tightly integrated, the network is can be measured by calculating the average of the shortest paths between each pair of two nodes in the network.<sup>39</sup>

<sup>39</sup> I do not mean "integrated" in the legal sense of "racial integration," though it is interesting to consider what the relationship might be between integrated in a graph theory, a sociological, and a legal sense. If two racial groups formed two completely separated cliques, it is hard to see how they would be integrated in any sense.

### C. Clustering

Real networks, unlike classical random networks, tend to be highly clustered. This appears to be a quite general feature of real networks of many different kinds, suggesting it is a feature of the mathematics of networks, something woven into their nature. A “cluster” (or “community”) in a network is a feature that is often obvious on inspection, as in the high school social network depicted above, with its racial and sexual clusters. This clumpiness of networks can be measured mathematically. Take, for example, social cliques. Suppose a person  $P$  has a group of friends. If every one of  $P$ 's friends is also a friend of every other friend of  $P$ , then  $P$  and his friends comprise, we say, a “fully connected” clique.



**Figure 7.** *Clique A* and *B* are “fully connected” within themselves, but only weakly connected to each other. The double arrows represent mutual friendships.

In *Figure 7* above, every member of *Clique A* is a friend of every other person in that clique, and similarly with *Clique B*. However, only one member of *Clique A* has a friend in *Clique B*, and vice versa. Both *Clique A* and *Clique B* are fully connected cliques. They are, however, only weakly connected to each other.

Network scientists use “clustering coefficients” to measure how integrated a network is. Clustering is a dimension network scientists take to be fundamental to the topology of any given network. The clustering coefficient is a measure, roughly speaking, of how tightly linked the clusters of a network are. The clustering coefficient of a particular node  $i$  is the ratio of the

actual number of links that connect to each other the nodes that are linked to  $i$ , to the number that there would be if  $i$  and the nodes it links to were fully connected. For example, if we treated the two cliques  $A$  and  $B$  above as separate networks (by snipping the link between them), the clustering coefficient of each node in *Figure 7* would be one. To get the clustering coefficient of an entire network, one takes the average of the clustering coefficient of every node.<sup>40</sup> In a legal citation network, if every case cited every other case in the network, its clustering coefficient would be one. If no case cited any other case, its clustering coefficient would be zero.

Clustering coefficients of networks can change over time. In their study of the scientific

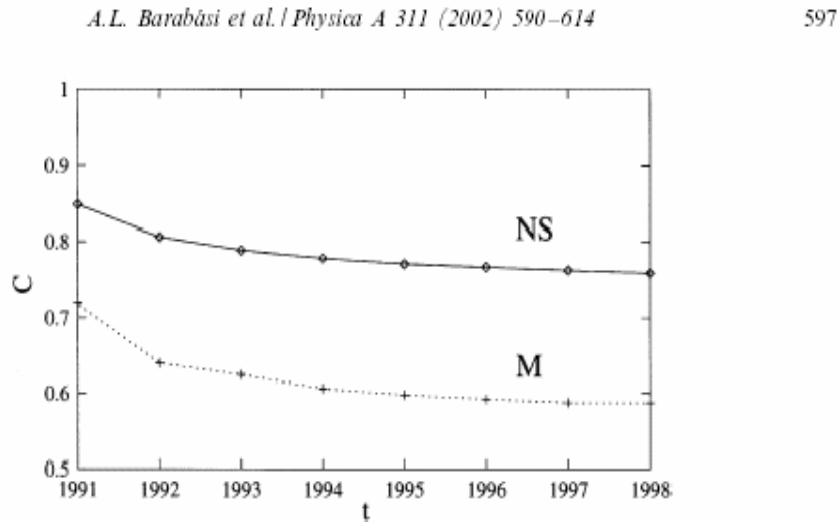



Fig. 4. Clustering coefficient of the M and NS database, determined for the cumulative data up to the year indicated on the  $t$ -axis.

collaboration network, Barabasi et al. also found that the clustering coefficient decreased over time for the period studied, and apparently approached a limit. Whether the scientific collaboration network would evince this effect over a longer period is unclear. However, the large, long-established and well-documented legal network would seem to be an excellent place to look for this phenomenon as well. If the same effect were found, it would suggest that the legal network was getting more integrated, in the sense of less clumpy (or, put differently, becoming one big clump), over time. If this were not found, it would suggest the opposite—that the legal network was organizing itself into clusters that were less and less tightly connected to one another.

<sup>40</sup> This can be a computationally very intense task, so one often makes do with a sample. See \_\_\_\_\_.



Clusters can be defined in different ways. We can see that it is natural to describe the network in *Figure 7* above as having two clusters. By snipping only one link, the network would fall into two pieces, and each would be fully connected within itself. Various mathematical algorithms can be used to define clusters more precisely than this intuitive illustration, but the intuition is sound. One useful algorithm involves the mathematics of flows through networks. Like many who live in Southern California, when I think of flows through networks, I think of traffic. *Figure 8* below shows San Diego and Los Angeles, connected by two freeways, the I-5 and the I-15. If one were to ask, how much traffic can flow between




**Figure 8.** Map of highways in Southern California shows a cluster at San Diego and a cluster at Los Angeles, weakly connected by two highways.

San Diego and Los Angeles, the answer would be, only what those two freeways can accommodate. The maximum flow is determined by the narrowest bottleneck in the system.<sup>41</sup> One way to define clusters in networks is to identify the “bottlenecks” and take those as the boundaries between clusters. The network of roads in southern California could be separated into at least two clusters, one around San Diego, and one around Los Angeles. In *Figure 7* above, the bottleneck would be the single link between *Clique A* and *Clique B*. We can reasonably say there are two clusters on either side of that link.


<sup>41</sup>

Clusters can be defined in different ways, as *Figure 9* below illustrates.



**Figure 9.** An example of different ways to define clusters. Clusters can be defined as the sub-networks within the circles, or as within and outside of the oval. From Jonathan A. McPherson, *Visualization and Interaction Techniques for Large Hierarchies and Graphs*, June 2004.

Depending on the algorithm, clusters can be defined as the sub-networks within the circles in *Figure 9*, or as those on either side of the oval. The coarser algorithm, which uses the oval, cuts only single links, and produces two clusters. The finer algorithm, the circles, cuts two or fewer links, and produces three clusters. The finer algorithm applies a stricter definition as to how tightly linked nodes must be to belong to the same cluster. Using finer and finer algorithms reveals progressively more tightly interconnected clusters. *Figure 10* shows the effects of applying a Markov algorithm repeatedly to a network with links of various weights,<sup>42</sup> illustrating how algorithms can reveal clustering at different levels of interconnectivity.



**Figure 10.** Markov algorithm applied sequentially to a network with links of varying weights reveals clusters. From Jonathan A. McPherson, *Visualization and Interaction Techniques for Large Hierarchies and Graphs*, June 2004

Depending on the network and the reasons for studying it, clusters can be defined in different ways and finding them can reveal different and sometimes surprising insights. In Part III.C below, I discuss how clustering analysis can be applied to the Web of Law.

<sup>42</sup> Markov algorithm

#### *D. Fitness and Network Evolution*

While Barabasi's scale-free model resembles the Web in overall structure, it is unrealistic in several important respects.<sup>43</sup> The model's mechanism of preferential attachment assures that the oldest nodes will have the most links. Yet in the Web and other scale-free networks, this is not true. There are many prominent exceptions to this rule, such as the website for Google, the most widely used Web search engine. Google is newer than many sites, yet far more popular than most. Some nodes apparently acquire links through some mechanism or mechanisms other than preferential attachment. These unusually attractive nodes are more "fit" than others, as some organisms in an ecological network are said to be fitter because they reproduce more. Google presumably possesses virtues that allow it to out-compete other nodes in the struggle to acquire links.

Network scientists are developing network evolution models that take different node fitnesses into account.<sup>44</sup> One of the more intriguing comes from Barabasi and [Giancomi] Bianconi. Bianconi discovered that the mathematical model physicists use to describe a quantum physical phenomenon called Bose-Einstein condensation ("BEC") also describes the evolution of networks in which nodes have different levels of fitness. BEC occurs in an atomic gas when it is cooled to only slightly above absolute zero.<sup>45</sup> When the gas is this cold, particles accumulate at a low energy level and (to use a layman's term)glom together, so much so, that the glob so formed can be described with a single quantum wave function. They become, in effect, an "it", in the form of a giant super-atom.<sup>46</sup>

---

<sup>43</sup> Another feature possibly prominent in the Web of Law is node aging, which does not occur in the BA model. I discuss aging networks at Part III.\_\_\_\_ *infra*.

<sup>44</sup> See Alain Barrat, Marc Barthélemy, Alessandro Vespignani, *Phys. Rev. Lett.* 92, 228701 (2004); S.H. Yook, H. Jeong and A.-L. Barabasi, *Y. Tu, Phys. Rev. Lett.* 86, 5835 (2001); Hyun-Joo Kim, Youngki Lee, Byungnam Kahng and In-mook Kim, *Journal of the Physical Society of Japan*, Vol. 71, No.9, pp. 2133-2136 (2002); Dafang Zheng, Steen Trimper, Bo Zheng, and P.M. Hui, *Phys. Rev. E* 67, 040102(R) (2003).

<sup>45</sup> Using lasers, physicists from the University of Colorado and MIT managed to do this, producing the predicted Bose-Einstein condensate and receiving the 2001 Nobel Prize in Physics for their trouble.

<sup>46</sup> This glob behaves in bizarre ways. Super-cooled helium that has undergone BEC behaves in a way to preserve its "single thingy-ness", very unlike what you would expect from a liquid, which in fact it is not - it's a Bose-Einstein condensate. If you have a bowl of helium BEC, and spill a little outside the bowl, it will defy gravity, and crawl back up the outside of the bowl to rejoin the rest of "it." If you spin the bowl of super-cooled helium, the bowl will spin, but the helium will remain still, with the bowl spinning around it. It is frictionless. The phenomenon is related to super-conductivity. See *SYNC*.

The mathematics that describes this process unexpectedly maps exactly the evolution of networks in which the nodes added have different levels of fitness.<sup>47</sup> Into the Bose gas model, one substitutes links for particles and nodes for energy levels. This model has an interesting consequence. In statistical mechanics, the behavior of a so-called Bose gas can be characterized completely by the distribution of the energy levels of the particles in it. So from observing the behavior of a gas, or its counterpart in this model, a network, one should be able to infer the distribution of energy levels, or fitnesses, of the nodes in the system. Thus, one should be able to characterize how much the distribution of links in the system (the “degree distribution”) is due to the rich getting richer (preferential attachment), and how much is due to the fit getting richer (more competitive nodes garnering more links). Assuming network scientists succeed in their efforts to develop the tools to measure this, it should be possible to measure whether nodes in a network are at the winning end of the power-law distribution because they are old, or because they are fit, or some combination. This would be a very useful property of a node in a network to be able to determine, whether the node was a web site, a legal case, or a law review article.

---

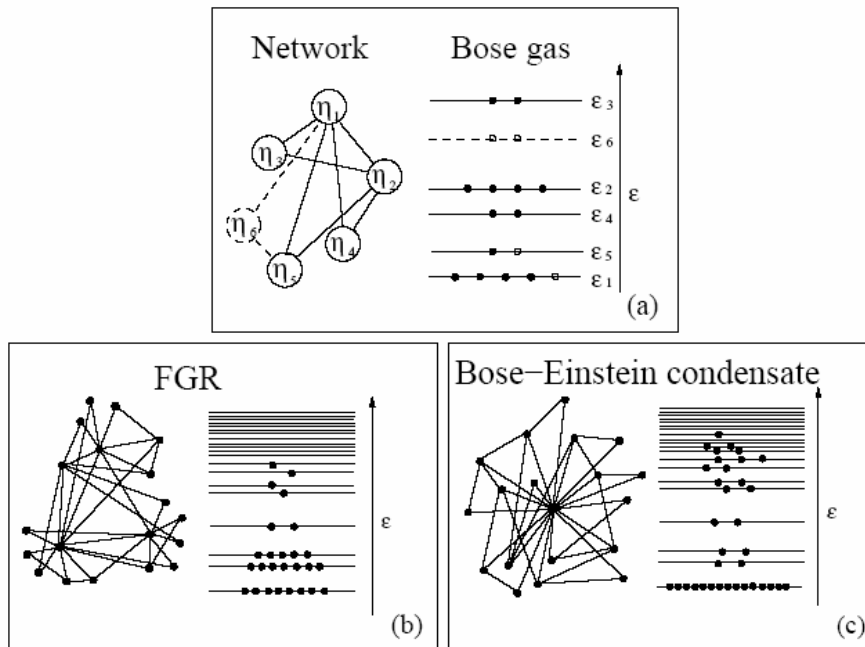
<sup>47</sup> If there is some deeper meaning behind this coincidence, scientists don’t know what it is. Coincidences of this sort are not unknown in science. See SYNC on pendulums.

**Figure 11. Mapping between the network model and a Bose gas.** From *A-L Barabasi, Emergence of scaling in complex networks*. Bornholdt and Schuster (Eds. at 77).

(a) On the left is a network of five nodes, each with a fitness given by  $\eta_i$ . An equation assigns an energy level  $\varepsilon$  to each fitness  $\eta_i$ . A link (or “edge”) from node  $i$  to node  $j$  corresponds to one particle at level  $\varepsilon$  and one at  $\varepsilon_j$ . The network evolves by adding a new node,  $\eta_i$  (dashed circle), which connects to 2 other nodes,  $\eta_i$  and  $\eta_j$ . These nodes were chosen by an equation that takes into account both preferential attachment and node fitness (or what I have analogized to stickiness). In the Bose gas model, this results in the adding of a new energy level,  $\varepsilon_i$  (shown as a dashed line), which is populated by 2 new particles (shown as open circles), and also in the deposition of 2 new particles to energy levels to which the new node is connected (i.e.,  $\varepsilon_i$  and  $\varepsilon_j$ ).

(b) In the “fit-get-rich” (FGR) phase of network evolution, there is a continuous distribution of “degree,” or number of links that a node has, with several high degree nodes linking the lower degree nodes together. In the energy diagram on the right, this shows up in fewer particles occupying higher energy levels.

(c) In the Bose-Einstein condensate, the fittest node (in the center, on the left side) attracts a finite fraction of all links. On the energy side, this corresponds to a highly populated ground level, and sparsely populated higher energy levels.



## II. THE WEB OF LAW

In Part I above, I introduced the concept of a scale-free network and explained some of the properties such networks have. In this section, I present evidence that the network of American case law is a scale-free network. In Part III, I suggest some of the consequences that the scale-free structure of the legal network implies, and indicate some promising avenues for future research.

### *A. The power-law degree distribution of scale-free networks*

One of the most striking features of a scale-free network is its “hub and spoke” structure. It was this property which led Barabasi and Albert to construct the scale-free network model in the first place. Their investigation of the structure of the Web revealed that the distribution of links to web pages did not fit a Poisson curve as predicted by the random graph model, but rather fell along a power-law distribution. This power-law distribution of degree (the number of links that a node has) is a fundamental feature of scale-free networks.<sup>48</sup>

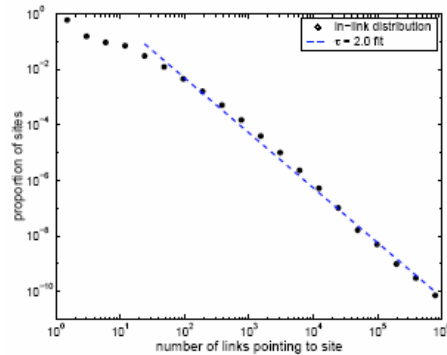

Many different networks share this power-law degree distribution property. *Figure 12* below shows the degree distributions of several different networks. *Figure 12(a)* shows the distribution of web page out-degree (number of hyperlinks running out of a web page) of a sample of the Web consisting of 200 million web pages. *Figure 12(b)* shows the distribution of in-degree in a sample of Web pages. *Figure 12(c)* shows the distribution of degree in the collaboration network of film actors. (In this network, actors are nodes, and a link exists between two actor-nodes if those two actors have collaborated in a film.) *Figure 12(d)* shows the distribution of the “language network.” In this language network,<sup>49</sup> words in a large text database are nodes, and two words are linked if they occur next to each other in some text. Each of the distributions above is displayed in the log-log format (each axis is on a logarithmic scale, with equal distances

---

<sup>48</sup> In Part III, \_\_\_ below, I discuss some respects in which the degree distribution departs from a power-law distribution.

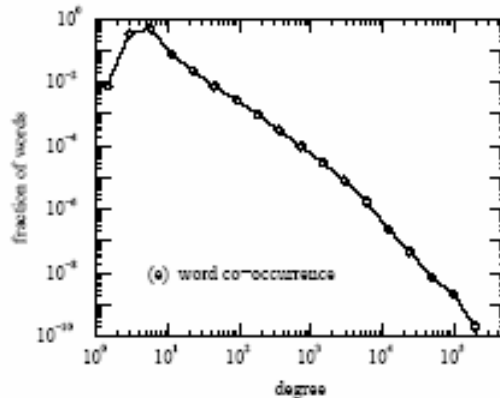
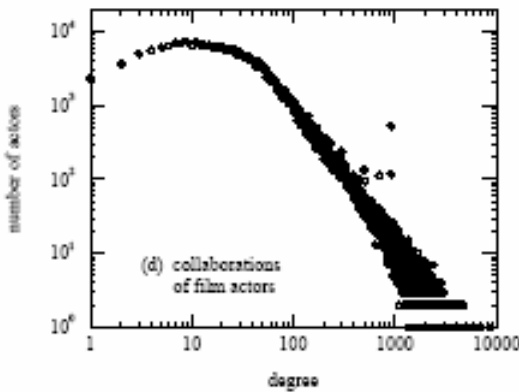
<sup>49</sup> There are different ways to view collections of texts as language networks. See

**Figure 12(a) (left).** Degree distribution of out-degree (links running out from web pages) from a portion of the web consisting of 200 million web pages.



**Figure 12(b) (right)** shows the distribution of in-links (links running into web pages) in a sample of the Web. Note that the  $y$  axis shows proportion of web sites. <http://www.hpl.hp.com/research/papers/weborder.pdf>

**Figure 12(c) (below left).** Degree distribution of links between film actors. Actors are nodes in the network; two actors are linked if they have collaborated on a film.



**Figure 12(d) (above right).** Degree distribution in the “language network.” In this network, words are nodes, and two words are linked if they occur next to each other in a library of text.

on the scale between 1 and 10 units, 10 and 100, 100 and 1000 and so on).<sup>50</sup> On a normal scale, the power-law distribution “hugs” the axes, so the log-log format is clearer.

In each of the networks displayed above, a few nodes (in the lower right quadrant) garner many links, while most nodes (in the upper left quadrant) garner relatively few links. This distribution is characteristic of a scale-free network, with its hub and spoke pattern. The

<sup>50</sup> Except for 12(b), the web in-degree. But making the  $y$  axis the proportion of web pages expressed as a power of ten, makes it equivalent to a logarithmic scale.

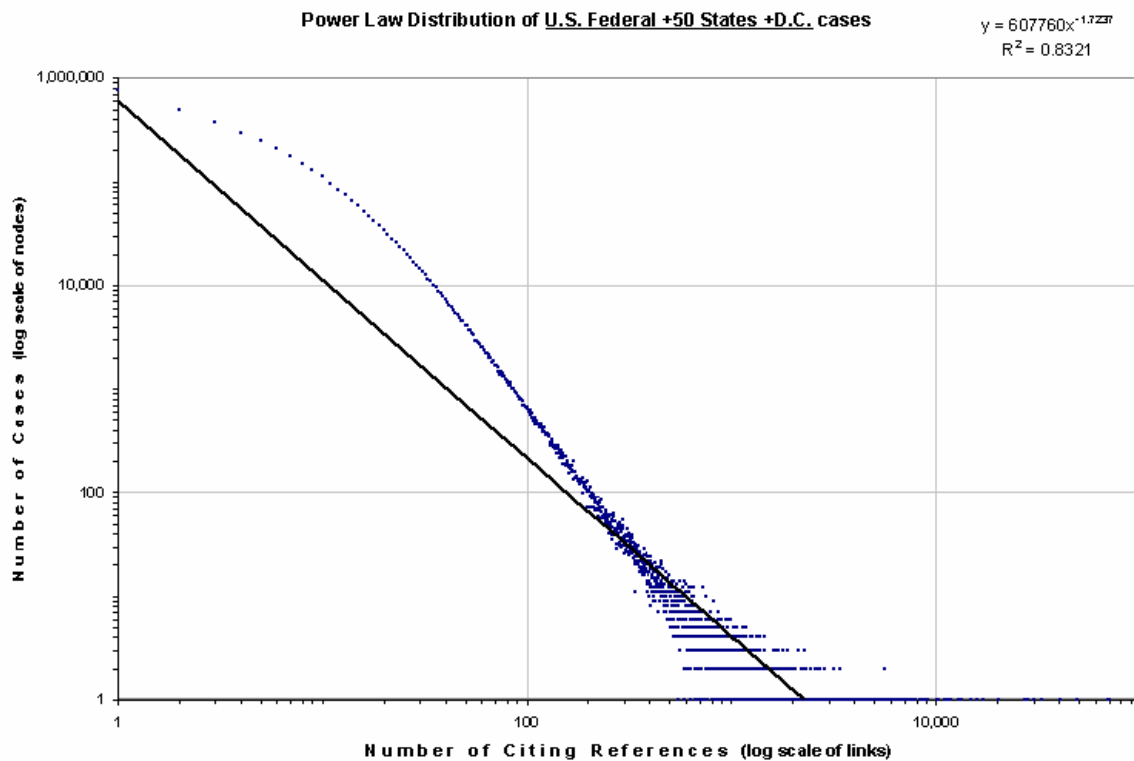
hubs are the nodes that have the majority of links, while the nodes at the ends of “spokes” have only a few links.

***B. The power-law distribution of legal citation frequency***

In reading Barabasi’s popular book *Linked*, I was struck by the similarity between his model for the mechanisms that generate the shape of the Web, and the processes by which common law systems grow. I reasoned that the “Web of Law” must be similar in shape to the Web. In the Web of Law, each case is a node, and it is linked to every case that it cites, and each of those cases is linked to the cases each of them cites, and so on. The legal network is a citation network. Citation networks have been studied extensively in the context of scientific papers, but legal citation networks have not been studied before now with the tools of network science, to my knowledge. Yet, the Barabasi-Albert model of scale-free networks, with its properties of growth and preferential attachment, seems especially well suited as a model of the legal network. In the BA model, nodes are more likely to get additional links depending on how often they have been linked to before. This is like the concept of legal authority. Cases that have been cited approvingly by judges in the past are seen as authoritative, and are more likely to be cited in the future, one would think. While sometimes newcomer cases gain authority quickly, the rich-get-richer or preferential attachment model of how cases accumulate legal authority (measured as number of citations) seems plausible, at least as a first approximation.

With this in mind, I contacted the LexisNexis corporation, the owner the well-known Shepard’s citation service. Every law student and lawyer knows how to “Shepardize” a case. The Shepard’s service publishes tables, both in hard copy and digital form, which show every



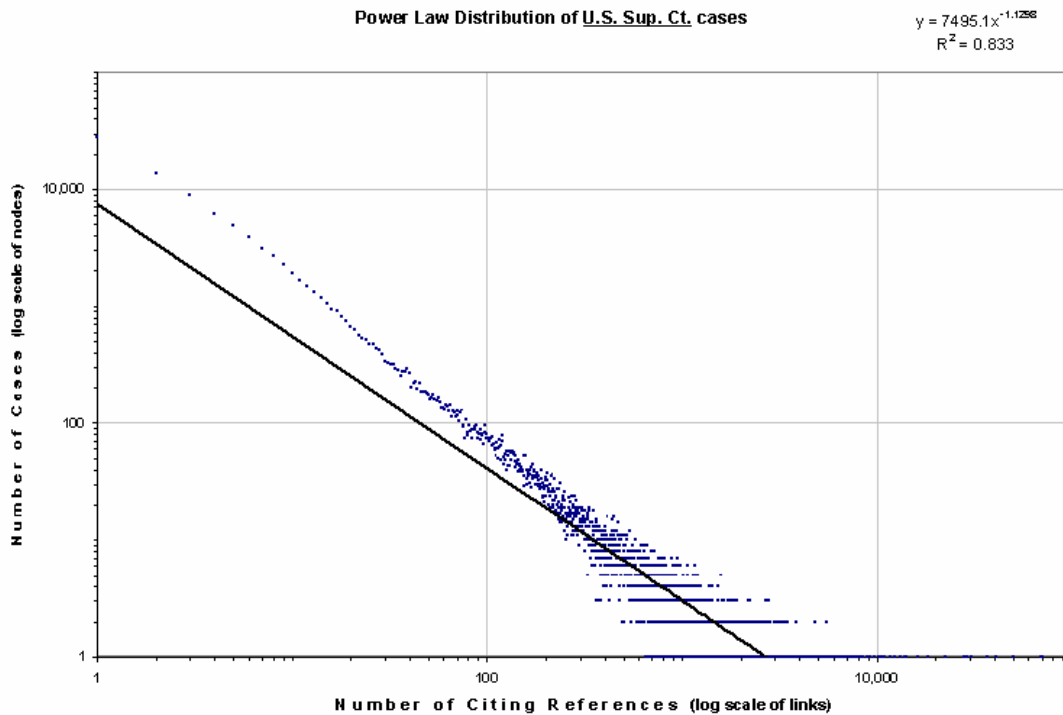


**Figure 13.** The power-law citation frequency distribution of American case law.

published decision in United States federal and state courts, and every case that cites such a published case. LexisNexis and Shepard's responded enthusiastically to my proposal that we undertake a study to determine the citation frequency (or "in-degree") of some subset of U.S. cases, to see whether the distribution conformed to a power law, as I strongly suspected it would, based on the BA model. In the event, economies of scale were such that it proved practical to determine the in-degree distribution for the entire Shepard's database of federal, state, and District of Columbia cases, consisting of nearly four million cases. While the study so far determines only in-degree distribution, which is an important, but fairly rudimentary dimension, it is, to my knowledge, the largest citation study of any sort, in terms of number of citations, ever performed, and certainly the largest legal citation study. The results are shown below in *Figure 13* (in the same log-log format used in *Figure 12* above).

*Figure 13* above conspicuously displays the same power-law degree distribution that is observed in other scale-free networks, such as those seen in *Figures 12(a)-(d)* above. The correlation between the actual distribution and a power-law distribution is not perfect. The R-

squared of .832, however, does strongly suggest that the mechanisms that produce a scale-free network are at work in the Web of Law.



**Figure 14.** The power-law distribution of citation frequency (in degree) of all U.S. Supreme Court cases. Courtesy of Lexis-Nexis/Shepard's.

The same citation-frequency (or in-degree) distribution analysis can be performed on subsets of the legal network.<sup>51</sup> Scale-free networks often have the property of self-similarity: Parts of the network have the same shape as larger parts of the network and as the whole network. So a subset of the legal network, if the network is scale-free, ought to have a similar sort of power-law distribution as the whole network. This appears to be the case for the Web of Law.<sup>52</sup>

Consider, for example, one of the most interesting subsets of U.S. cases for legal scholars, the set of U.S. Supreme Court cases. The citation frequency distribution of all U.S. Supreme Court cases

<sup>51</sup> When the data of this study is released to me by LexisNexis, I anticipate having a program that will be able to quickly determine the degree distribution of virtually any jurisdictional subset of cases in the Shepard's database.

<sup>52</sup> See *infra* at \_\_\_\_\_.

is shown in *Figure 14*. Supreme Court cases conspicuously display the same power-law distribution of frequency that is observable in the legal network as a whole, and also in more familiar scale-free networks, such as the Web, that are shown in *Figure 12*.<sup>53</sup> (I should note, however, the distributions for Supreme Court cases and for all U.S. cases also appear to differ from each other in that the former is less concave to the powerlaw trend line. I conjecture this is both significant and interesting, and discuss a possible explanation below in Part III.\_\_\_\_.)

The distribution of citation frequency of the legal network indicates that the Web of Law is a scale-free network, in all likelihood. This is a preliminary result, but it is significant, because legal scholars have not hitherto been aware that the legal network has this structure. Further analysis needs to be done to determine those dimensions of the Web of Law which network scientists typically analyze, such as the exponent of the degree distribution of the legal network, the extent to which degree conforms to a power-law distribution. However, merely from knowing that the Web of Law is a scale-free network, many interesting insights can be inferred and many possibly fruitful avenues for future research identified. In the next Part, I explore some consequences of the scale-free structure of the legal network.<sup>54</sup>

### III. CONSEQUENCES OF THE STRUCTURE OF THE WEB OF LAW

The scale-free structure of the Web of Law is not merely a curiosity. Lawyers, judges, and law professors have long resorted to metaphors of webs, trees, and bramble bushes to evoke the structure of law.<sup>55</sup> These are just metaphors, but the metaphors are attempts to get at

---

<sup>53</sup> Note that the R-squared of the whole legal network frequency distribution to a power-law distribution is .832, and that of USSC case citation frequency distribution has an R-squared of .833.

<sup>54</sup> Attentive readers will notice that while the distribution of the citation frequency in the Web of Law and the in the other networks illustrated in *Figures 12(a)-(d)* shows evidence of power-law distribution, the fit is hardly perfect. Most of the network graphs show a downward bending hook or hump at their northwest most end, and the graph for the whole Web of Law in particular shows a conspicuous bulging, which is concave to the line indicating the power-law distribution in log-log format. They also often show a truncation of the power-law distribution at the other end as well. Why is this? These effects are probably caused by other network effects that are largely beyond the scope of this introductory treatment, but which should certainly be investigated. Most important, I conjecture, is node aging. It may be that as cases get older, they begin to lose their ability to attract new links, preferential attachment notwithstanding. See L. A. N. Amaral, A. Scala, M. Barthelemy, and H. E. Stanley, *Classes of small-world networks*, <http://amaral.northwestern.edu/Publications/Papers/Amaral-2000-Proc.Natl.Acad.Sci.U.S.A.-97-11149.pdf>.

I discuss this further below at Part III.\_\_\_\_.

<sup>55</sup>

what we can now describe with far more precision as the network structure of law. The Web of Law shares a mathematical structure with many other evolving networks. This structure is being studied intensely by network scientists, and is gradually yielding up its secrets. The scale-free structure of legal authority implies that law has an overall shape and internal structure that can be profitably studied in new ways.

The legal network grows as judges write opinions and cite cases and other authorities in those opinions. Similarly, scholars add to the legal network by writing articles and treatises in which they cite other authorities, and which are cited by other authorities. A great deal of information is embedded in these citations. Judges cite the cases that they think are the most relevant to the cases they are deciding. When two judges deciding different cases cite some of the same authorities, it is a clue that those cases are relevant to each other. By way of millions of decisions regarding what cases or other authorities to cite, the legal network organizes itself into what network scientists call “clusters” or “communities.” In other scale-free networks, these clusters form not just structures in link topology, but also structures in semantic topology. This implies a significant result for general jurisprudence. Scale-free networks such as the Web and, I argue, the Web of Law, spontaneously organize themselves semantically, that is, according to meaning or topic, where topics are defined not by some external authority, but as a property emerging from millions of informed decisions. This in turn implies that our common law system, and indeed, any common law system, should have an organic organization that can be mapped and studied, potentially shedding great light on the internal structure and intellectual history of law. Interestingly, this is a consequence of the mathematics of any legal system with rules and principles of the kind that causes it to generate a scale-free network, as ours does.

This network perspective should also change the way we look at legal precedents and such doctrines as *stare decisis*. As I discuss in Part III.\_\_\_\_ below, the respects in which the Web of Law departs from the pure Barabasi-Albert model are important and interesting. They suggest that legal precedents may have a natural life span, a term over which they grow in authority, attain a kind of maturity, and then decline into relative obscurity. Scientific authorities do this in scientific citation networks, and the distribution of degree in the legal network gives at least the impression of doing something similar.<sup>56</sup>

---

<sup>56</sup> At this stage, this is a conjecture on my part, based merely on the inspection of the distribution of the data and looking at other models in the literature. However, at this early stage of applying network

This Part is divided into eight subparts, in roughly increasing interest and importance. In subpart A below, I discuss whether the legal network is a “small world,” and what the answer to that question might suggest about the legal system. In subpart B, I discuss clusters in networks and how this concept applies to the Web of Law. In subpart C, I explain how the structure of the legal citation network supports the claim that law is intellectually coherent. In subpart D, I consider the significance of “hub cases,” cases that have far more citations than most. In subpart E, I discuss the idea of node fitness and how it can be applied to the legal network. In subpart F, I explain how the network structure of the Web of Law might be exploited to improve computerized searching in legal networks. In subpart G, I consider how network theory can be applied to our concepts of authority and precedent and the related doctrine of *stare decisis*. I criticize a model of change in statutory interpretation proposed by Professor Dan Faber, which relies for empirical support on his observation that the citation frequency of a sample of textual interpretation cases he examined are power law distributed. In part H, I propose the hypothesis that most American cases have a limited life span of authority, and that this life span varies by the prestige of the court producing the decision. This hypothesis needs to be confirmed by further research, but the similarity of case law and scientific paper citation frequency distributions, and the contrast between the distributions of all U.S. cases to all U.S. Supreme Court cases, strongly suggest this is the case. Finally, there is a brief conclusion.

#### *A. Is the Web of Law a “Small World”?*

Counting federal and state cases, the American legal network has about four million nodes.<sup>57</sup> We could imagine, as a first approximation, a network in which these nodes were arranged randomly, like buttons thrown on a ballroom floor. Each node could be linked only to the two nodes closest to it. If this is what the Web of Law looked like, any two nodes chosen randomly would probably be very far apart, in terms of the number of links one would have to traverse to get from one to the other. However, in real, non-random networks, short cuts between nodes often greatly reduce the average distance between nodes. In the widely studied Watts-Strogatz model,<sup>58</sup> for example, as illustrated in *Figure 14* below, random rewiring of a

---

theory to law, I think it is useful to make plausible conjectures that can be confirmed or falsified by further research.

<sup>57</sup> If one included law review articles and statutes, it would be even bigger.

<sup>58</sup> Watts, *supra* note \_\_\_\_\_.

Smith:

circular network results in the rapid decrease of the average distance between nodes. Many real world networks, from groups of insects, to human social groups, to the Web, are “small worlds” in this sense. They have links that connect otherwise widely separated nodes, which greatly shorten the shortest path between any two nodes.




Fig. 3. Average separation in the M and NS databases. The separation is computed on the cumulative data up to the indicated year. The error bars indicate the standard deviation of the distances between all pairs of nodes.

Measuring the legal network in this dimension to see how “small” it is, is important because the “diameter” of a network gives an indication of how well integrated and interconnected that network is. The possible implications of discovering that the network of American law is a small, or not so small, world, are many. If one could connect any state decision with a decision in any other state, for example, in only a few links, this would suggest the various state legal systems were not as widely separated as one might otherwise think. Similarly with the decisions of federal and state courts. Comparing the diameter of the state and federal case law networks as a whole, to the diameters of each set of cases, federal and state, considered respectively as separate networks, would give one an empirical measure of how separate from or integrated with each other these nominally separate systems actually were. This could confirm notions we already have about federalism, or lead to unexpected insights.

Networks can also become more or less integrated over time. Barabasi and \_\_\_\_\_ found, for example, that in two networks of scientific collaborators, the average separation between nodes was decreasing over time, and was approaching a limit. In these networks, scientists are nodes, and they are linked if they have collaborated together on a paper. Would the average separation in the legal network behave in the same way? That is, would one find that over time, different areas of law were growing less separated from one

another? If one found that the state and federal legal networks were not growing more integrated over time, or that integration was approaching a limit, this would suggest that there was something about the combination of a common law and a federal legal system that caused the initially separate legal systems to resist complete integration over time, and that this was a property that was woven deeply into federal (in the sense of multi-level) common law systems. On the other hand, it is conceivable that the data would indicate the separate legal systems in our republican form of government are gradually becoming more integrated, with no limit in sight. If this were the case, it would suggest the federal legal system would not “naturally” maintain itself, but rather, left to its own devices, the law of the states and of the federal government would gradually merge into a single, integrated system. Finally, this might turn out to be true in some areas of law, but not in others.

Legal specialization seems, based on casual empiricism, to be increasing constantly or even accelerating. If average node separation is increasing over time in the legal network, as it might be doing, this would suggest that as the legal network gets bigger, it gets more spread out—that, unlike scientific collaboration networks, it does not become more integrated over time, but less. If this is the case, we might expect the insularity of different areas of law to increase, with no limit in sight. If analysis of the data supports this view, it would be an important result in the sociology of law. On the other hand, actual research might show that this conjecture is false. Claims of inevitable specialization and growing isolation of different areas of law might be overstated. It might turn out to be that law manages to integrate itself, notwithstanding its ever growing size. If this is true, it would also be an important conclusion. Either way, the investigation is well worth pursuing. It is worth noting that in terms of size, long time period covered, and quality of documentation, the legal network presents an excellent opportunity to network scientists to study network evolution. It would not be surprising if the Web of Law was the largest and best documented citation network in existence.

### ***B. Legal clusters***

Scale-free networks tend to share topological traits. This implies a deep fact about all common law systems, and indeed all legal systems, that grow over time and have a doctrine of precedent and authority that creates a preferential attachment mechanism. These features (which may be combined with other mechanisms as well) generate a scale-free network of legal authorities. This scale-free network, in turn, is probably organized in clusters. Because the



Web of Law is a scale-free network, it probably consists of clusters of cases which are relatively tightly linked within themselves, but more loosely linked to each other, analogously to the loosely linked social cliques and highway networks illustrated above, and to the structure of the World Wide Web. Furthermore, these clusters probably correlate highly with underlying legal semantics. That is, cases in the same legal cluster are likely to be related to each other semantically, in terms of meaning and subject matter.<sup>59</sup>

In terms of the “micro-motives” that produce the “macro-behavior,”<sup>60</sup> linked legal cases ought to be, one might think, even more closely related semantically than linked web pages. In deciding a case on, say, an issue related to the contract law doctrine of the parole evidence rule, a court is likely to cite cases that are as relevant to the parole evidence rule as it can. The court is also likely to cite cases which are jurisdictionally relevant. The judge will prefer to cite a case from his own court or from a higher court in its jurisdiction, than from some remote jurisdiction. These citations will be produced by persons who are presumably intimately familiar with the case at hand and with the relevant law. Judges would thus seem to have stronger motivations and greater expertise for choosing relevant links than would Webmasters. The link topology to semantic topology congruence, therefore, would seem to be even tighter in the legal network than in the Web, where this phenomenon was first described.<sup>61</sup> On the other hand, the standards by which relevance is judged in law are much stricter than in the Web. In any event, a great deal of local expertise is embedded in the citations judges choose to make in writing opinions, more than any central authority, such as indexers or treatise writers, is likely to have. Looking at the law as a network may allow us to exploit the wisdom of judicial crowds.<sup>62</sup>

If link topology maps well onto semantic topology in the Web of Law, then analysis of clustering in the legal network would give us a picture of the natural organization of law. Law presents daunting problems of organization. It seems probable that a mode of organization that is naturalistic, that is, an organization that is *found* in the legal system, rather than imposed upon it, is likely to yield insights about law. One may analogize this to biological taxonomy. Biological organisms evolved in a certain way and order. Taxonomies based on evolution will

---

<sup>59</sup> Gibson, Kleinberg, Raghavan, *Inferring Web Communities from Link Topology*, <http://citeseer.ist.psu.edu/cache/papers/cs/596/http:zSzzSzwwww.sims.berkeley.eduuzSzcourseszSzis296a-3zSzf98zSzreadingszSzgibson.pdf/gibson98inferring.pdf>

<sup>60</sup>

<sup>61</sup> *See supra* note \_\_\_\_.

<sup>62</sup> wisdom of crowds book

probably be more useful to scientists, and more productive of additional insights, than would be artificial taxonomies based on more superficial features, such as whether an organism is sea-dwelling or land-dwelling.<sup>63</sup> The natural topology of law might or might not conform well to conventional organizations of law.

### *C. Is Law Intellectually Coherent?*

The data presented in this article is relevant to the global issue of the intellectual coherence of law. It suggests that, in terms of intellectual coherence, the legal network could look a lot worse. Consider, for example, the radical position that the entire legal system is intellectually incoherent. Perhaps no legal scholar actually espouses this extreme position, but we might take it to be an exaggerated version of what some Critical Legal Scholars have claimed from time to time.<sup>64</sup> What would such a claim really mean? In terms of the concepts used in this article, one way to model an extremely intellectually incoherent legal network would be as a decision making procedure that produced a citation network that was a random graph. In a legal network for which a random graph was a good model, judges would literally choose which cases to cite at random, closing their eyes and opening the reporters to find cases to cite, the way some people use the Bible for oracular guidance. Even if they did not do this, judges might do something such that they might as well be randomly choosing cases to cite, for all that their practice was driven by some coherent set of legal ideas. So if their choosing which cases to cite were guided by some intellectually garbled practice, they would choose cases in a way that might as well be random, for all that some other judge could reproduce the first judge's choices, in a similar case.

However, if this were what judges were doing, citation patterns would *not* look like they in fact do. Far from being random, the legal citation network appears to have the same basic structure as do citation networks of scientific literature, which presumably set a high standard for intellectual coherence.<sup>65</sup> This does not mean, of course, that law is a science. It does suggest, however, that there are similarities between the structure of law and the structure of science, at least in terms of what might be called the structure of authority. What we know of the legal network so far is at least *consistent* with the claim that it is highly intellectually coherent. The legal network does not look remotely like a random graph.


---

<sup>63</sup> <http://www.ohiou.edu/phylocode/>

<sup>64</sup> most ridiculous crit available.

<sup>65</sup> Scientific citation studies

**Figure 15.** Visualization of a scientific citation network containing approximately 34,000 nodes. Clustering is clearly evident. The legal citation network would be about 100 times larger. McPherson, *Visualization and Interaction Techniques for Large Hierarchies and Graphs*, June 2004. <http://wwwcsif.cs.ucdavis.edu/~mcpherso/largegraph.html>



#### *D. Hub Cases*

Scale-free networks are as integrated as they are because of the relatively few nodes that have many links. In social networks, some people know far more people than most.<sup>66</sup> Some of us are such people, but most of us are not. These widely-connected people are hubs, through whom go the links that tie together people who would otherwise be widely separated. Some of these hub people can be seen in the densest parts of the high school social graph above. If the legal network is well integrated, it is probably because particular cases, and cases decided by particular courts, are the hubs that make the legal network smaller, and more integrated, than it would otherwise be.<sup>67</sup>

What are the hubs in the Web of Law? Powerful courts, such as the U.S. Supreme Court, and the Federal circuit courts, have probably decided many of the hub cases. Hub cases might also include many procedural decisions, because procedural jurisprudence applies to many different substantive types of cases.<sup>68</sup> Similarly, some constitutional cases might prove to be hubs, because the principles they articulate apply to many cases in otherwise diverse legal areas. It also might be that particular judges are responsible for more than their share of hub cases. In any event, there would be a sense, not to be exaggerated, but still important, in which these hub cases help hold the legal system together. Discovering and describing the set of hub cases in the Web of Law would obviously be worth doing.<sup>69</sup>

It would be useful to see how legal scholarship fits into the legal network as well. Would law review articles be as, more, or less clustered, than judicial decisions, for example?<sup>70</sup> Do legal scholars, or at least some legal scholars, act as “integrators” between or among different areas of law, like the centrally located persons in the high school social diagram? Or do they tend to remain clustered and insulated in their specialized areas? Some important law review articles or other legal scholarship would probably appear as hubs, more at the center

---

<sup>66</sup>

<sup>67</sup> I do not mean to suggest that the citation network is the only legal network. There is obviously a legal social network, consisting of lawyers who know each other professionally. There are also networks of persons in legal authority, in the sense of persons who can legally exercise power over other persons. To name a few. The citation network probably contains the most information, however, about the content and evolution of legal doctrines and principles, which makes it of great interest to legal scholars.

<sup>68</sup> Faber at \_\_\_\_\_. Faber observes that many of the most frequently cited cases in his data are procedural cases. I suspect many of the hub cases are procedural cases.

<sup>69</sup> Future research.

<sup>70</sup> Faber notes that the most frequently cited cases by judges in his data have little in common with the cases most frequently cited in law review articles. Farber at \_\_\_\_\_.

than the periphery of the Web of Law. I would suppose that the legal scholarship network and the network of case law and statutory authority would interpenetrate, but not overlap completely. Certain important cases and statutes, such as *Roe v. Wade*<sup>71</sup> and RICO<sup>72</sup>, would probably sport coronas of law review articles that cite them, and articles that in turn cite them, and so on. Other cases would enjoy no such entourage.

In a scale-free network, some nodes are hubs, but most are not. In the legal citation network, hub cases are cases that are cited much more frequently than the vast majority of cases. One must be cautious, however, in interpreting what this means. The most frequently cited case in American case law, with about 17,000 citations to its credit, is the *Liberty Lobby*<sup>73</sup> case, a pillar of the U.S. Supreme Court's jurisprudence of summary judgment. It is, without doubt, an important case. Presumably it is cited so often because the federal courts, and state courts following federal courts' lead, handle many motions for summary judgment, and very often when they do, they resort to the law in *Liberty Lobby*. It does not follow, however, that it is therefore the "most important" or "most authoritative" case in American law. On the other hand, it would be hard to deny that *Liberty Lobby* provides law related to the core of what federal courts do.

Would a set of hub cases, defined in some appropriate way, do a good job of capturing the fundamental doctrines of American law? Or, perhaps one should ask, would the so-called "fundamental" doctrines of American law do a good job of capturing the hub cases? Obviously, questions like these stir some deep jurisprudential waters. A legal realist might say that what the law actually *is* has a great deal to do with what judges actually *do*, and therefore might be inclined to give some significant weight to citation practices.<sup>74</sup> A legal formalist, on the other hand, might believe law had some abstract structure that sustained itself independently of judicial citation practice.<sup>75</sup> I do not mean to espouse any particular jurisprudential view in claiming that the legal network is an object worthy of study, and I acknowledge that the relationship between citation practice and legal importance is complex. Whatever one thinks is the ultimate jurisprudential significance of the Web of Law, there is no denying it is the repository of an enormous amount of information about law, information that can be extracted by the right tools.

---

<sup>71</sup> cite

<sup>72</sup> cite

<sup>73</sup> cite

<sup>74</sup> Leiter on legal realism.

<sup>75</sup> Legal formalism

### *E. Fitness in the Web of Law*

Nodes in scale-free networks arrange themselves along a power law curve, in terms of “in-degree,” or how many links each node has, as *Figures 12(a)-(d)* and *13* above illustrate. In the log-log format used above, the nodes with the most links, the hub nodes, are those to the southeast or lower right, while the much more numerous nodes with few or no links, are those to the northwest, or upper left. The single most northwesterly point in *Figure 13* represents a group of roughly 400,000 cases that have been cited only once, [by themselves<sup>76</sup>]. While not mathematically accurate, one can fancifully picture some 400,000 nodes, each a neglected case, stacked on top of each other at that northwesterly-most point, in a veritable catacomb of forgotten case law.<sup>77</sup>

The nature of a scale-free network, however, is to evolve as new nodes and new links get added. Measuring the in-degree of a network, such as the Web of Law, is like taking a snapshot of it at a particular time. Can we picture more dynamically what this evolution of new nodes and links looks like? To describe this evolution, it would be helpful to have a metaphor that was more evocative for law professors than an atomic gas being chilled to the brink of Bose-Einstein condensation. Coming up with a better metaphor takes some doing. The following rather elaborate fable is my attempt to convey the evolution of a network like the Web of Law, which has nodes of varying fitness. I ask the reader’s indulgence for its fanciful quality. It has the advantage of being, though less rigorously accurate, at least more vivid for most people than analogies from quantum physics. So, let us imagine that when a court issues an opinion, it produces a small bound copy of it, which weighs a few ounces. To each opinion is attached a small helium balloon, like children sometimes get at parties. Lighter-than-air balloons will be our metaphor for links, which in the Web of Law are citations by other cases. Cases start out with only one citation, the case’s citation of itself. One balloon does not have enough lift to pull a case off the ground, however. As of mid-2004, about 400,000 cases, some old, and some new, were in this state of obscurity. In our fable, these 400,000 booklets, with only one balloon attached to each, are scattered across a large meadow, earth-bound. Yet other cases have been cited more often. Some (such as *Liberty Lobby*) have been cited *much* more often, and so have many more balloons attached to them.<sup>78</sup> As cases get more citations, and thus more balloons get attached to them, they rise progressively farther off the ground.

<sup>76</sup> Every case treated as citing itself, so no cases with zero cites (?)

<sup>77</sup> Catacombs in Lima, Peru.

<sup>78</sup> (perhaps associates do the tedious work of attaching the balloons to the correct cases).

Because the atmosphere gradually thins, making the balloons steadily lose lift, the cases array themselves at different altitudes, depending on how many balloons (or citations) are attached to them. All the cases weigh the same, so (let us assume) the altitude of each case is determined just by how many balloons are attached to it. This presents a curious spectacle. Many cases, about a tenth of the total, just sit on the ground. Most of the rest are at different altitudes, but most are still pretty close to the ground. A relatively few cases have a lot of balloons attached to them. The number of cases one sees falls rapidly as one gains altitude, and a relative few soar high above the others.

If we were to film this evolving spectacle and then study the film, we would observe a number of things. First, we would see that some cases worked their way to a high altitude apparently by virtue, at least in part, of having been around longer than most others. Age has its advantages. These cases would get balloons tied to them relatively slowly, but surely, and so rise in the ranks, some making it into the highest rank. Other cases would get a few balloons attached to them, but then stall somewhere in the middle distance. Still others would rise quite quickly, getting balloons attached to them at a rapid pace, and rise anywhere from somewhat to much faster than other cases that were making more stately upward progress. And many cases would just sit there on the ground.

In this informal model, balloons get attached to cases partly as a matter of preferential attachment, that is, as a function of how many balloons a case has attached to it already (how “rich” it is in citations). We can imagine that judges looking for cases to cite look up and see floating cases, and their attention is drawn to those that are floating highest and have the most balloons attached to them. Those attention-grabbing cases are the ones that are most likely to get still more balloons attached to them. But this is not the only dynamic at work. Cases also get additional balloons attached to them as a matter of how “fit” they are. Just from watching the film, we cannot tell what it is that makes some cases more fit than others. However, in the Barabasi-Albert preferential attachment model, it should be possible, given that we know how many links a node has already, to calculate what the probability is that a given node will get an additional link (or balloon), when a set of new links is added to this network. Thus, we can imagine that every week, a new set of cases (with single balloons attached to them) is brought out to the meadow, and new balloons added to those in the air, or on the ground, according to which cases (besides themselves) those new cases cite. If we know what the probability is that a particular case will get a new balloon attached to it, then we should be able to determine which

cases, over an appropriate number of cycles, are getting new balloons attached to them at a rate *greater* than that predicted by the preferential attachment model. How much greater that rate is, would be a measure of the “fitness” of the case. “Fitness,” in this sense, is thus a residual category. It is that portion of a node’s ability to garner links that is not explained by its already having a certain number of links. When a case garners citations beyond what the preferential attachment model would predict, we say the case is relatively fit. How fit is a matter of how many links it gets beyond what preferential attachment predicts.

The above model, though informal, does offer, I believe<sup>79</sup>, a method that could be implemented mathematically to measure the “fitness” of nodes in the Web of Law, such as cases and law review articles.<sup>80</sup> Why should one want to do this? First, once one understands that the Web of Law is a scale-free network, and realizes the power of the preferential attachment mechanism, it is natural to wonder how much the ongoing authority of particular cases and other legal sources is due merely to the fact that their authority is so well established. The scale free structure of law is evidence of the intuition that at least some authority in law is to a significant extent self-justifying or at least self-reinforcing. One might say that the mechanism of preferential attachment is an explanation, at a certain level, of the ability of the legal status quo to sustain itself. Legal authority does not just sustain itself; It accumulates. At least, that is what its scale free structure suggests. I do not mean to suggest that the mathematics of networks is the deepest available explanation for the dynamics of legal authority, only that this dynamic, whatever its ultimate causes are, evidences itself in the structure of the legal citation network. Presumably other factors, social, political, economic and legal, explain the preference that reveals itself in preferential attachment. The scale free structure of legal authority, however, indicates this preference, whatever causes it, is a reality that requires an explanation.

Be that as it may, cases that possess authority<sup>81</sup> that cannot be explained by preferential attachment are evidence of other dynamics at work. Some cases might appear to be cited more than others because they are better reasoned, more persuasive, or more accurately interpret statutory language or legislative intent. If we actually measured the fitness of these cases, we might find, however, that the citation frequency of some of them was accounted for by

---

<sup>79</sup> A better model would also take account of the complications I raise below in part III. \_\_\_\_\_.

<sup>80</sup> Many specification would have to be added to the model, of course. Such as how many cases were added each week, month and year. This growth should presumably be based on the actual growth of the legal network.

<sup>81</sup> Or lack authority, as aging nodes do. *See infra* at \_\_\_\_\_.



preferential attachment. Or, we might find that they were actually cited *less* than one would expect, given their age, like a retirement savings account that would be impressive for a 20 year old, but not for a 65 year old. This in turn might allow us to detect early on that some cases that were not yet authoritative, were heading in that direction, and conversely, that authoritative cases were in the process of losing their authority. While the task would be computationally intense, one can imagine in principle a search engine with the option of searching for “emerging authorities” that would calculate a fitness index for cases within given relevance parameters (produced by a text-based search, for example) and would highlight cases of high fitness, or rank them in fitness order. (Other possible improvements for legal search engines based on the Web of Law concept are discussed in subpart D below.)<sup>82</sup>

Measuring the fitness of cases and other legal authorities could also be useful to legal historians. While one should be cautious not to exaggerate the power of quantitative methods in history, the uniquely well-documented Web of Law would seem to present new opportunities for legal historiography. In principle, one should be able to measure the changing fitnesses of cases over time. Perhaps, for example, one could see the waning of the authority of older U.S. Supreme Court cases after legal watersheds such as the New Deal. Historians might be able to actually see when formerly important cases began to be overruled *sub silentio*. Conversely, one might see be able to see other cases spring into prominence. As the English legal historian S. F. C. Milsom demonstrated in his magisterial *Historical Foundations of the Common Law*,<sup>83</sup> law is an extremely opportunistic discipline, with precedents and doctrines being exploited by advocates and jurists to solve immediately pressing problems in ways impossible to predict *ex ante*. Here legal evolution seems closely analogous to biological evolution. In biology, one sees a similar sort of opportunism, with organisms adapting organs to almost unbelievably “ingenious” uses. Working backward, one can sometimes reconstruct when and sometimes why the adaptation was “made” by natural selection. It may not be too much to hope that in the patterns of citations, mundane though

---

<sup>82</sup> Subject to lots of qualifications; trends could reverse themselves etc.

<sup>83</sup> For example, the first case I worked on as a law clerk for the late Judge George MacKinnon of the U.S. Court of Appeals for the D.C. Circuit involved a garden variety wrongful termination suit by a fired employee who claimed he had an oral employment contract for a term of five years, the enforcement of which the employer said was banned by the statute of frauds. A clerk for another judge on the appellate panel claimed the trend in the relevant state law was “clearly” not to bar claims with the statute of frauds. In fact, it was not clear at all what the “trend” was, and the case ended up being rather bitterly disputed among the judges. It would be interesting and useful to have a methodology that could detect the emergence of trends in legal authority on more than just an “eyeball” basis. Whether courts would use it or not, is another question, of course.

they might seem, one could see, as in fossils buried in layers of stone, similar branchings in the evolution of law.

Law professors would probably also like to know how “fit” their scholarship was, compared to that of other scholars. I have heard anecdotes to the effect that when Professor N was being considered for tenure at University X, an opponent of the candidate’s did a citation count of the candidate’s articles and found they had been cited only  $R$  times, a number implied to be shockingly low. It is a fair guess that the method used in such attacks was scandalously unscientific. But citation frequency and other patterns can divulge large amounts of information, if used properly. By (to put it crudely) backing out preferential attachment, we can get a clearer idea of how “fit” a piece of scholarship is. Types of influence could also be studied. For example, an article might have little impact on other scholars, but a major impact on judges, or vice versa.

#### ***F. Exploiting the scale free structure of the Web of Law to improve legal research***

If the Web of Law has the same basic structure as the Web, it follows that searching for relevant authorities in the legal network could be improved by drawing on techniques from search strategies that work well on the Web. Searching in scale-free networks is now a highly developed and technical field. Offering a superior search algorithm, Google has become a multi-billion dollar public company. Compared to technology routinely used to search the Web for recipes and movie schedules, the technology used to ferret out relevant precedents in the Web of Law on such weighty matters as corporate takeovers and the death penalty seems relatively primitive.

The insight that the legal network and the Web are very similar in structure suggests a new approach to searching for relevant legal authorities might improve, perhaps greatly, the quality of search engines used on electronic legal databases. As a casual experiment, the reader may wish to compare, just in terms of a subjective assessment of the quality of results, a search conducted on Google, and a search conducted on Lexis or Westlaw. I usually find that a Google search gives me, within the top five or so results, one or more web pages that are highly relevant to what I am looking for. I rarely get such subjective satisfaction from searches on Lexis or Westlaw, where I expect to have to wade through dozens of results to find one that is

highly relevant to my query. This, of course, is just a subjective report of experience, but I suspect many legal scholars have similar experience.

Why, to put it bluntly, does Google do so much better than Lexis or Westlaw at producing relevant responses to queries? It may be the difference in quality is partly illusory. Law may be so much more precisely defined and categorized than more general information on the Web that we judge the relevance of legal search results by a much higher standard. However, it may well be there is more to the difference than that. Google exploits the information imbedded in the linkages of web pages to rank the relevance of search results. Google uses its patented PageRank algorithm to rank search results produced by a more conventional text-based system.<sup>84</sup> Roughly speaking, the Google search engine first produces a set of results based on the occurrence of terms on a web page. Then those web pages are scored for relevance on the basis of how many web pages link to them. This score is calculated so that being linked to by a webpage that itself has many links increases a web page's score by more than being linked to a webpage with only a few links. So if a web page is linked to by Yahoo or CNN.com, it will have a higher relevance score than a page linked to by, say, some obscure blog.<sup>85</sup>

It may seem rather mysterious that exploiting linkage information in this way gives Google such a powerful tool for assessing "relevance." After all, if one is looking for a web page with a good recipe for brownies, why should a web page to which Yahoo happens to have linked be judged a superior source of information on chocolate confections? It helps to recall that all of judgments of relevance in this context are strictly matters of probability, which is not the way lawyers usually think of relevance. Yahoo is heavily linked to in part because Web surfers have judged it to be a good source of information on what are websites with good information. If Yahoo links to a site, it is probable that that site has good information about whatever it happens to be about. Many different Web masters endorse Yahoo's authority by linking to it, and Yahoo confers that authority on a brownie recipe by linking to it. This is hardly foolproof guarantee, and the network science on authority and relevance is far more sophisticated than this exposition might suggest. As a general matter, however, it is intuitive that

---

<sup>84</sup> Anatomy of a search engine

<sup>85</sup> For a fascinating analysis of the influence of blogs, and exploration of the implications of their power law distribution of links, see Daniel W. Drezner & Henry Farrell, *The Power and Politics of Blogs* (Working Paper, August 2004), <http://www.danieldrezner.com/research/blogpaperfinal.pdf>.

sites that many other web pages have found to be good sources of information, are more likely than other web pages to link to sites that have good information about the topics they address.

In the Web of Law, analogously, a great deal of information is embedded in the links that courts have made among cases, and which legal authorities in general, including statutes and law review articles, have made among one another. As a first approximation of an improved search engine, search results produced by the Boolean and other text-based systems, such as those employed by Lexis and Westlaw, could be improved by allowing users, at their option, to display results in the order of citation frequency.<sup>86</sup> The search results produced by a Boolean search, such as, say, “fiduciary w/10 duty and derivative and directors” (to take a broad corporate law issue search one could conduct on Lexis) ought to themselves fall along approximately a power-law distribution, in terms of citation frequency.

This result is mathematically almost trivial. A random sample taken from a population with a given distribution will probably have the same distribution as the population. Even if mathematically trivial, however, it is practically interesting. It indicates that for any set of cases (or other legal authorities) determined by some text-based algorithm (such as a Boolean search) to be relevant, there will be a relative few cases that have been cited many times and a relatively large number that have been cited rarely. For the same reason that highly linked web pages are more likely to be relevant to a query than rarely linked web pages, frequently cited cases are more likely to be relevant to a query than rarely cited cases. To extend the analogy to Google, one could rank search result cases by a score that gave greater weight to being cited by cases that were themselves frequently cited, which would probably improve the relevance of search results a great deal.

These Google-ish approaches, however, are probably not ideally adapted to a legal citation network. While the Web of Law and the Web are both scale-free networks, they are different in respects important in judging relevance. The Google approach depends to some extent on a fact about the Web that is not true of the Web of Law. While some web pages exist solely to archive static information, such as, say, a website that reproduces the books of the Bible, many web pages exist to display information that is constantly changing and growing. Alternatively, if a relatively successful web page does not offer changing content, it at least offers

---

<sup>86</sup> In fact, fastcase.com does this. It also allows one to rank results based on citation frequency by the cases the text based search produces. See [www.fastcase.com](http://www.fastcase.com)

content of such value that as the Web grows, new links are constantly being added to it. This is much less so with the legal citation network. Because legal doctrines change over time as courts and legislatures make new law, it is quite possible for a case with many citations to be anything but authoritative on any given question of law, because a more recent case from a higher court may have overruled it. In a related way, cases sometimes fall out of favor even though they have never been expressly overruled. Thus, if a case has been cited five hundred times, but not once in the last forty years, it is probably less authoritative, for the purposes of a lawyer, than a case that has been cited fifty times in the last two years. Also, the legal system is more hierarchical than the Web, at least in some senses. If a case has been cited a hundred times by various federal district courts, that probably gives it less weight as authority than being cited twenty times by the U.S. Supreme Court. The Web of Law I am discussing is a citation network, while the Web is more a living electronic library designed so anyone can add to the collection. A citation is something that occurred in the past, even possibly the remote past, while a hyperlink is something that must be maintained to be an ongoing part of the Web. For these reasons, if one were to use a version of Google's PageRank to judge relevance in the Web of Law, the results might depart markedly from what lawyers, judges and legal scholars would judge to be legally relevant. Such an approach may not adequately account for how law is constantly moving away from its past, nor would it account for the hierarchies of legal authority.

Rather than a purely Google-type approach, therefore, an approach more specifically adapted to the legal network would be preferable. The first step is to account for the fact that a link or citation in the legal network, unlike a link in the Web, inevitably decays over time. The law is more like a tree than an electronic network. In an electronic network, all parts have to be activated constantly by the flow of electrons. In the Web of Law, the surface, like the outmost layer of a tree, is where the life is. Past decisions provide a kind of skeleton that supports the contemporary life of the law, but current law is vital in a way that law of the relatively remote past is not. The most straightforward way to operationalize this fact about legal authority is simply to apply a discount rate to the addition to the relevance score that a case would get from a particular citation. The further in the past was a citation, the less it would add to the relevance score of a case. So, a 1960 case that was cited once in 1970, would be scored as less relevant than a 1960 case that was cited once in 1980, other things being equal. Similarly, the weight added to relevance by citation should be scored by the level of the court that did the citing. So, a 1960 case cited once in 1970 by a federal circuit court would rank as less relevant than a 1960 case cited once in 1970 by the U.S. Supreme Court. These ranking rules could be

applied recursively. Thus, a case cited by one 1990 Supreme Court decision would rank higher than another case cited by one 1980 Supreme Court decision, all other things being equal. All other things would not be equal, of course. In practice, one would calculate the weight being cited by a particular case would add to a case a number of levels back. So the relevance of a case would be determined by the weights of the cases that cited it, and those weights would be determined by the weights of the cases that cited them, and so on. In principle, one could go through the entire Web of Law this way, back to the beginning of the Republic, but in practice, this would not be necessary. By moving ten or so steps back in the legal network, one would probably glean enough information to order search results in a way that conformed better than existing methods to actual human lawyers' judgments of relevance.<sup>87</sup>

A search engine that ordered results in terms of cases (or other authorities) that were the most "fit" would be useful in some applications. Legal scholars and legal practitioners, for example, often want to know what are the cases, articles, or other authorities that are attracting the most attention lately. Recent cases may not have had time to accumulate many citations, and yet they might still be attracting a relatively large amount of attention for their age. These might be called "emerging authorities." Emerging authorities can be identified by calculating the fitnesses of cases returned in Boolean or other text-based search. Doing this would require calculating an expected citation weight for cases of a particular type and age, then comparing that weight to the weight of cases produced by the text-based search. Type of case could be identified with a general area with which a lawyer wanted to remain abreast. More generally, the fitness of each case produced in a text-based search could be calculated, and then cases could be ordered according to fitness, or alternatively, according to some weighted average of fitness and relevance, as defined above.

---

<sup>87</sup> Legal relevance also differs from Web relevance because of the uniquely legal idea of jurisdiction. To a lawyer practicing in the Ninth Circuit, a Ninth Circuit case is much more relevant than a case decided by the Second Circuit, and vice versa. Any scoring system should be adaptable to jurisdiction. Thus the searcher, for example, should be able to order, at his option, the Ninth or the Second Circuit as the more relevant jurisdiction. This election would affect the scoring of relevance recursively. Thus, cases cited by Ninth Circuit decisions would add more to the relevance score, when the Ninth Circuit was elected as the leading jurisdiction, than would cases cited by Second Circuit decisions. Ideally, searchers would be able to indicate the order of relevance of jurisdictions. For example, a searcher might indicate the Ninth Circuit was the most relevant jurisdiction, then circuits that were adjacent to the Ninth Circuit, and so on.

### *G. Faber's Tectonic Model of Statutory Interpretation*

In a forthcoming article, Professor Daniel Farber argues that the power-law distribution of citations in certain sets of cases is consistent with what he terms a “tectonic” model of statutory interpretation.<sup>88</sup> The reference comes from the Gutenberg-Richter power law, which describes the distribution of earthquake magnitudes. There are many earthquakes of small magnitude and a few of large magnitude, and the distribution of earthquakes rather surprisingly fits a power-law distribution. The mechanism that produces this distribution is thought by some to be phenomenon of self-organized criticality, which is also evidenced by power-law distributions.<sup>89</sup> Farber argues that power-law distributions in statutory interpretation cases indicates or may indicate that a kind of rupturing, rather than gradual process, characterizes how legal interpretations of important texts, such as statutes and constitutions, change over time.<sup>90</sup>

In this section, I argue briefly that Farber's thesis, while it may be correct, is not supported by the data he presents. Neither is his data inconsistent with his theory. In fact, as *Figure 14* above shows, all U.S. Supreme Court cases are power law distributed. It would be quite difficult to come up with any subset of Supreme Court cases that was *not* power law distributed. It may well be true that shifts in constitutional and statutory interpretation usually occur, as a historical matter, more often as ruptures or paradigm shifts than as a result of a gradual process. However, the fact that the citation frequencies of the entire Web of Law are power law distributed, and more particularly the in-degree of Supreme Court cases appears to be even more uniformly power law distributed than all U.S. cases, implies that the evidence Farber adduces does not in fact support his tectonic model any more than it supports a completely gradualist model.

In his study, Farber selects cases that involve textual interpretation.<sup>91</sup> He then measures the frequency of citation of these cases and finds that those citations are power-law distributed. However, this result does not show what Farber argues it does. Because the entire set of

---

<sup>88</sup>Daniel A. Farber, *Earthquakes and Tremors in Statutory Interpretation: An Empirical Study of the Dynamics of Interpretation*. Minnesota Law Review, 2004 <http://ssrn.com/abstract=519883>

<sup>89</sup> For a good popular account of the Gutenberg-Richter power-law distribution and possible explanations of it, see MARK BUCHANAN, *UBIQUITY: WHY CATASTROPHES HAPPEN* 43-62 (2000).

<sup>90</sup> The idea that statutory interpretation occurs by rupture instead of gradually is attributed by Farber to William N. Eskridge, Jr., *Dynamic Statutory Interpretation*, 135 U. PA. L. REV. 1479 (1987).

<sup>91</sup> Farber at 13.

American case law is power-law distributed, and is, I argue, a scale-free network, any (approximately) random subset of the entire set will also be power-law distributed. Thus, one could choose the Supreme Court's contract cases, or all American cases in which the name of one of the parties begins with the letter "R," and one would find their citation frequency, or "in degree," was approximately power law distributed. (This generalization is subject to an important qualification below.<sup>92</sup>) Some readers will no doubt find this assertion counter-intuitive, so an analogy may be helpful. Most readers are no doubt familiar with the ubiquitous normal distribution, or bell curve. Normal distributions characterize many phenomena, the magnitude of which is independently influenced by many different factors.<sup>93</sup> For example, the height of adult American males is normally distributed around a mean of about five feet, nine inches tall. If one were to measure the height of every adult male who was in attendance one Sunday at a San Diego Padres baseball game, one would find that those heights were normally distributed around some mean. Now suppose one selected out some subset of those men according to some arbitrary criterion, such as those that had the letter R in their California Driver's License number. One would also find that the heights of the men in this subset was normally distributed. Similarly, if one pulled out Republican men, Democratic men, men with red hair, men wearing no socks, and so forth. It is possible that the mean would vary slightly, especially if one selected a subset by some criterion that was relevant to height. So, for example, Asian American men tend to be shorter than the overall American average. However, *all* of the groups would be normally distributed with respect to height, even if the mean varied somewhat.

It is the same with power-law distributions. Virtually any subset of cases pulled out of the Web of Law as a whole will be power law distributed, or more precisely, as approximately power law distributed as the whole Web of Law is. Therefore, the fact that textual or statutory interpretation cases are power-law distributed is not evidence of anything in particular about those cases, except, as I claim here, that they are part of a much larger, scale-free network. Cases that are not about textual interpretation are power law distributed. Cases about any topic *X* are power law distributed.

Farber also claims that the power-law distribution of citation frequency is evidence of a sudden, tectonic process in the changing of interpretations, rather than a gradual one. A

---

<sup>92</sup>

<sup>93</sup> Smith, partnership equality

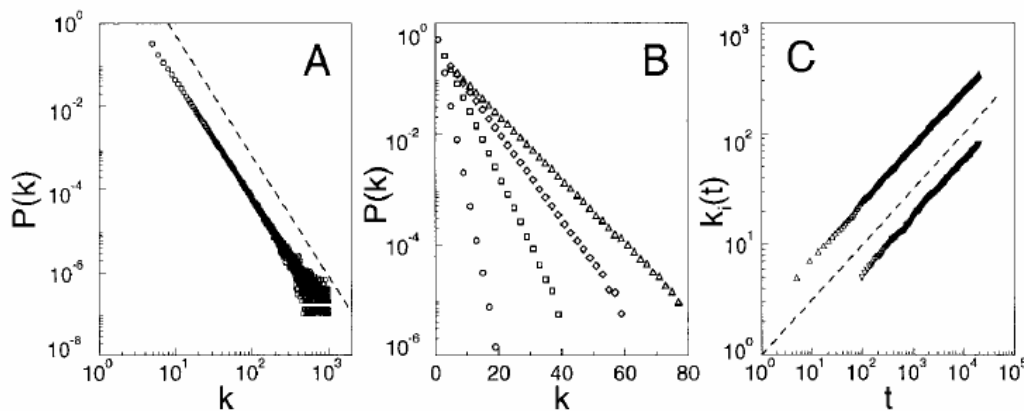


gradual process, he suggests, would produce a normal distribution of case citation frequency. However, this is not true either. Growth of a network over time and preferential attachment would account for the power-law distributions of citations. Thus, a legal system that was *entirely* based on preferential attachment, and was thus utterly traditionalist and hide-bound, and never changed its interpretation of anything, so that the *only* factor determining the probability of a case's getting additional citations would be how many citations it already had, would also, as we have seen, produce a power-law distribution of citation frequency. In fact, certain departures from this utterly precedent-bound approach, would cause *departures* from a power-law distribution. In order to get something like a normal distribution of citation frequency, one would need to cite cases at random, at least with respect to previous degree of citation. That is, in order to see a Poisson distribution, which is not normal, but is roughly bell shaped, how often a case had been cited in the past would have to not predict whether it would be cited again as the network of law grew. This is not really of matter of sudden or gradual change.

Network theory does suggest, however, a way to test empirically whether some cases or interpretative approaches acquire authority suddenly. It may turn out, and indeed I suspect it is true that, as Farber hypothesizes, shifts in interpretations of important constitutional clauses occur suddenly rather than gradually. Network theory offers a way to test this hypothesis. Assuming we equate authority with citation frequency, which is probably at least roughly the case, it should be possible to measure the rate at which cases acquire citations, and compare that to the rate they would have done so, just by virtue of preferential attachment. Some cases might be seen to take off dramatically in terms of in-degree, analogously to the way in which the Google web site garnered links far faster than a mere preferential attachment mechanism could explain. Paradigm-shifting cases would appear, therefore, in my model, as cases of extraordinarily high fitness, as that concept is explained above. If Farber's hypothesis is true, some appropriately identified set of cases about constitutional interpretation should have significantly higher fitnesses than does a suitable set of control cases (such as all Supreme Court cases). A similar approach could be used to measure the fitness of legal scholarship, such as law review articles. I believe that it is likely that certain cases and law review articles would appear to have much greater fitness than most of their kind and would qualify as paradigm-shifting authorities, but only actually measuring fitness would reveal this for certain.

### H. Node Aging and the Life Span of Legal Authority

Recall that when Barabasi and Albert first began investigating the structure of the Web, they expected to find that links were distributed according to a Poisson distribution, as classical random graph theory would predict.<sup>94</sup> Instead, they found a power-law distribution of links, and they developed their model of network growth and preferential attachment to explain that unexpected distribution of link degree. The BA model stimulated a huge amount of research into network structure by scientists trained mainly in the mathematical techniques of physics, to the point that now the statistical mechanics of networks can be seen as a sub-discipline of physics in its own right. It soon became clear that as important an advance as the BA model was, it presented a highly idealized picture of the Web and other real networks.



**Fig. 2.** (A) The power-law connectivity distribution at  $t = 150,000$  ( $\circ$ ) and  $t = 200,000$  ( $\square$ ) as obtained from the model, using  $m_0 = m = 5$ . The slope of the dashed line is  $\gamma = 2.9$ . (B) The exponential connectivity distribution for model A, in the case of  $m_0 = m = 1$  ( $\circ$ ),  $m_0 = m = 3$  ( $\square$ ),  $m_0 = m = 5$  ( $\diamond$ ), and  $m_0 = m = 7$  ( $\triangle$ ). (C) Time evolution of the connectivity for two vertices added to the system at  $t_1 = 5$  and  $t_2 = 95$ . The dashed line has slope 0.5.

sciencemag.org SCIENCE VOL 286 15 OCTOBER 1999


!

**Figure XX.** A shows the degree distribution of a network generated by the BA model.

This can be demonstrated by comparing a graph of the degree distribution produced by a numerical simulation of the BA model, with a graph of the degree distribution of a real network. *Figure XX* above shows the degree distribution that was produced by such a numerical simulation, which applied the network growth and preferential attachment mechanisms of the BA model. Note how closely the distribution in *Figure XX.A* above conforms to the power-law distribution trend line. Real networks, however, frequently depart

<sup>94</sup> See *supra* at \_\_\_\_\_.

from this idealized, numerical simulation. For a dramatic example, consider the network of film actors, displayed above in *Figure 12*, and reproduced again below in *Figure 1X* to facilitate



*Figure 1X.* Network of Film Actor collaborations, with region of exponential decay (1) highlighted.

comparison. This network, like many other real networks, does not conform at its tails to the pure power-law distribution produced by the BA model. This deviation is of particular interest to us because this effect is prominent in the scientific citation network, and, as referring back to *Figure 13* will remind the reader, is quite apparent in the Web of Law as well. Interestingly, however, this exponential decay is *not* as apparent in *Figure 14*, the graph of citation frequency for U.S. Supreme Court cases. For convenience of comparison, the degree distribution of the Web of Law, of U.S. Supreme Court cases, and of the citation network of a large sample of scientific papers, is shown below in *Figures 1Y A, B, and C*. It is apparent how the Web of Law graph resembles that of the degree distribution of scientific papers, in that in the northwest region, the distribution curves downward, departing from a power-law distribution and turning into (what looks like) an exponential decay.




Figure X3.A *Scientific Paper Citation Network*. From S. Redner, How Popular is Your Paper, Eur. Phys. 1998

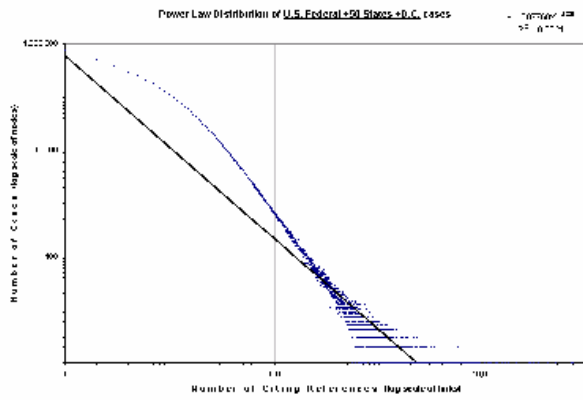


Figure X3.B. *Web of Law*. Distribution of citation frequency of all U.S. state and federal cases




Figure X3.C. *U.S. Supreme Court Cases*. Distribution of citation frequency of all Supreme Court cases.




FIG. 1. Data on the network formed by scientific publications (nodes) and citations (directed links). Upper panel, circles: The number of papers published in a given year from 1987 to 1998. Triangles: The total number of citations made in papers published in 1998 and referring to papers published in a given year [25]. The data for both curves have been extracted from the ISI database [29]. Lower panel: The average number of citations (incoming links) a paper received in 1998 as a function of the paper's publication year. The values are obtained as the ratio between the values of the two curves in the upper panel. Considering only papers more than 3 years old (published before 1995) the rate of obtaining new citations decreases with age. This indicates that aging is an important feature of citation networks.

they get.

A similar phenomenon probably explains the citation frequency distribution in the Web of Law as well. The power of cases to draw citations probably decays with age, and this effect is probably most marked with cases that will never be cited a lot in any event. Some preliminary citation counts I have done by hand suggest this effect may be so strong in the least prestigious courts that their citation distributions do not significantly fit a power-law distribution, and so are not accurately characterized as scale-free networks. Put another way, most of the decisions emanating from some courts might be found in the northwesterly region of the Web of Law. This in turn suggests that decisions of some courts may age very quickly indeed. What time scales we are talking about must be determined by further research, like that which has already been done on scientific paper networks.

What explains this distribution in the network of scientific paper citations? This is an interesting question for legal scholars, for the implication of the similarity of the distribution of legal citations to that of scientific paper citations is that the mechanism may be the same, since it produces such a similar effect. The probable explanation, which needs to be confirmed by further research, is that that scientific papers and legal precedents, especially those that are cited only a few times, age quickly. Scientific papers and cases lose their power to attract new citations over time. This fact can be seen directly in data concerning the distribution of citation frequency by publication date of scientific papers, shown in *Figure 1Z*. These data show that scientific papers reach their prime, in terms of attracting cites, when they are about three years old. Papers older than that are on average cited less, and still less, the older they get.

By contrast, consider the distribution of citation frequency of cases decided by the U.S. Supreme Court. As *Figure IX.C* shows, it demarks a much straighter line, very similar in form to that produced by the numerical simulation produced by the BA model. While some slight concavity to the trend line may be noticed, it is markedly less than that apparent for the Web of Law as a whole, or for the scientific paper network, or for the network of those notoriously quickly aging film actors. This is a remarkable contrast. It suggests, as is consistent with our intuitions, that Supreme Court precedents age much less quickly than do those cases that are less frequently cited in the Web of Law as a whole. It also suggests that some courts may work under what is effectively a different regime of *stare decisis* than do other courts. Further research should reveal what the average life span of precedents produced by the U.S. Supreme Court and other courts are.

Figure IXX below<sup>95</sup> graphs numerical simulations of networks in which the nodes age (lose their ability to attract new links) at different rates. This simulation illustrates vividly how various magnitudes of node aging truncates the distribution of citation frequency. In this preliminary article, my object at this point is mainly to observe that the distribution of the Web of Law appears merely by inspection to resemble more closely a network with aging nodes than one without aging nodes, and conversely that the network of Supreme Court authorities looks relatively more like a network with nodes that age more slowly. Obviously, more research has to be done here, and very interesting research it will be.




FIG. 3. Distribution of the connectivities for several values of the aging exponent: 1)  $\alpha = 0$ , 2)  $\alpha = 0.25$ , 3)  $\alpha = 0.5$ , 4)  $\alpha = 0.75$ . The inset shows  $\log P(k)$  vs.  $k$  for  $\alpha = 1, 5$ , and 10. Note that, in the later case, all three curves are nearly the same.

Part of the reason this research will be interesting is that node aging has a profound influence on the overall structure of the network, and that node aging seems to involve critical values. So in *Figure IXX* we see, for example, that when the aging exponent gets to equal 1, then the distribution acquires a shape that further increases in aging does not affect. So the action, so to speak,

<sup>95</sup> S.N. Dorogovtsev and J.F.F. Mendes, *Evolution of reference networks with aging*

occurs when aging is between two critical values. More interestingly, node aging has a profound effect on the structure of the network. In the pure BA model, where there is no node aging, when a scale-free network gets large enough, clustering disappears. The network becomes one big cluster, so to speak. This is where things start to get rather strange, but fascinating nevertheless. This suggests that in a legal system where precedents really did stay precedents, and did not grow outmoded with the passage of time, all cases would come to be equally closely connected to each other. Law would become unified in a way someone might wish for, but certainly seems alien to current understanding, where different areas of law address different topics. Alternatively, consider a regime in which cases gathered *much more* authority as they got older. Perhaps my Federalist Society friends will forgive me if I label this, the Originalist regime. In this case, the citation pattern undergoes something like a Bose-Einstein condensation, with cases surrounding a few central hubs. *Figure 1* below shows the effects of negative aging, no aging, and aging on a simple circular network. Where the aging exponent




FIG. 1. Change of the network structure with increase of the aging exponent  $\alpha$ . The aging is proportional to  $\tau^{-\alpha}$ , where  $\tau$  is the age of the site. The network grows clockwise starting from the site below on the left. Each time one new site with one link is added.

is negative 10, (the “Originalism” regime), nodes surround one old, authoritative node. Contrast this with the network with very rapid aging, in which the aging exponent is 10. Here the network also has a very simple structure, the only authority being the last case decided. This is interesting for many reasons. To name one, it suggests there is a fundamental tension between some version of Originalism, which could be interpreted to hold that the closer a precedent was to some founding moment, the more authoritative it was, and the hope that there would be multiple authorities in a legal network. The more cases gain prestige for being older, the more the network wants to be centralized. But on

the other extreme, cases that age too quickly produce no structure of authority at all. This is obviously over-reading the metaphor of a simple model, but it is suggestive, nonetheless.

Where in this array of possible networks does the real Web of Law fall? Where would one want it to fall? These are very abstract questions, but it does no harm to ponder them. If research reveals that cases in the Web of Law age quickly, and that all but a few Supreme Court cases are temporary creations of courts that will cease to exert much authority after a few years, that would certainly shed a new light on the enterprise of law. It would make much more plausible the claim that, globally speaking, courts are really just administrative bodies that settle disputes as they arise, and their pretences of doing so consistently with long-standing law is ill founded at best. On the other hand, if some significant portion of the Web of Law ages slowly, then this would suggest, as a very general matter, that the Rule of Law is not merely a pious fiction. To able to approach such a grand question empirically, even if not to settle it definitively, seems worthwhile work indeed. It may well be that the Web of Law is somewhere between these two extremes. This would, in a sense, be the most interesting result of all. For it seems that networks whose nodes age, but not too quickly, generate the most interesting networks of all. They do not evolve into one big cluster, nor do they evolve an equally boring structure, a network with no memory, in which the only structure is the most recent decision. They evolve into highly clustered, hierarchical networks.

If it turns out, as it well might, that the life span of legal authorities in common law systems is a magnitude that contributes to, or rather, is necessary, to that system's evolving a highly complex structure, that would be a remarkable thing. It may be in that case that it would be as true to say that the law organizing itself drove the aging of cases, as it would be to say the aging of cases caused law to organize itself. It may be that law has to be complex—have many fields, and sub-fields, and sub-sub-fields, and many different authoritative voices—to perform the tasks we demand that it perform in a complex society. Or perhaps, and perhaps not inconsistently, law just wants to be organized.

#### IV. CONCLUSION

American case law and other legal authorities are organized in a certain way, as a web or network, with its nodes connected by the links of citations. This network can be considered as a mathematical object whose topology can be analyzed using the tools pioneered by physicists and others who wanted to explore the structure of the Web and other real networks. The Web of Law has a structure very similar to that of other real networks, such as the Web and the network of scientific papers. The Web of Law is in substantial part a scale-free network, organized with hub cases that have many citations and the vast majority of cases,



which have very few. The distribution of citation frequency approximates a power-law distribution, as is common with real scale-free networks, with truncations at either extreme of its distribution, which is also common.

Many promising hypotheses can be generated by considering the law as a scale-free network. State and federal systems can be examined empirically to measure how well integrated each is with itself, and with each other, and how this is changing over time. Legal authorities can be measured to determine whether their authority is emerging or declining. Institutional bodies, such as courts, can be examined in the same way. Clusters of cases, which will reveal the semantic topology of law, can be mapped to determine whether traditional legal categories are accurate or require reform. These methods can be operationalized in computer programs to improve the efficiency of searching electronic legal databases. The topology of American law can be compared to that of other legal systems to determine whether legal systems share universal architectural features, and in what respects different systems are unique. Changing dynamics of the citation frequency and the fitness of particular cases can be studied over historical periods to test historiographical hypotheses. So, for example, Farber's hypothesis that changes in constitutional interpretation occur suddenly, and many others, may be tested rigorously. The dynamics of authority in law generally can be studied much more rigorously. The mere fact that law is a scale free, not a random network, suggests a high degree of intellectual coherence, contrary to what some critics of the American legal system have suggested. The shape of the degree distribution graph of the Web of Law, in its similarity to the scientific citation network, also suggests that cases age, in the sense of losing the ability to attract citations, over time, just as scientific papers do. Yet Supreme Court cases seem to age more slowly. What cases are most permanently authoritative, and what cases are the least, is worth investigating. How nodes age profoundly affects overall network structure and therefore affects the shape of the Web of Law. Network theory hints at complex, but analyzable, interactions between the legal doctrines of precedent, and the systems of common law and multiple sovereignties.

Because law grows and because it has doctrines of authority, it creates a network of a certain shape. The sort of legal system we have has features that cause law to spontaneously organize itself. This is the product of laws that govern networks of computers as inexorably as they govern networks of cases, laws arising from the underlying mathematics of networks, which are indifferent to the substance that instantiates them. Yet in ways we do not yet understand

well, these laws enable a legal system to self-organize and function, perhaps far better than a system without these features could do.

Understandably enough, network scientists have been far more interested in the network of scientific publications than they have been in the Web of Law. Part of the purpose of this article is to advocate a collaboration between legal scholars and network scientists to explore what is probably the largest and best documented citation network ever created. The Web of Law is probably about four times as large as the largest scientific citation network, and stretches back some two centuries. And this is just the American legal system. As any legal scholar who has waded through the Uniform System of Citation (the “Blue Book”) knows, the network of legal citations is documented with a thoroughness and precision that borders on fanaticism, which (however painful to learn) is a feature highly desirable in a network to be analyzed. Legal scholars do not have Ph.D.’s in statistical mechanics, and most physicists have only the vaguest idea of how law really works, but collaboration can overcome those problems. Because of the pioneering work of companies such as LexisNexis, legal databases are huge, precisely documented, and readily accessible. They present a perfect opportunity for the application of network science. This research would produce new knowledge of general jurisprudence that has simply been impossible until now, when we have the necessary advances in network science, the fast computers, and the existence of a complete record of the legal network in electronic form, waiting to be explored.