

How Long is the Coastline of the Law?
Thoughts on the Fractal Nature of Legal Systems

David G. Post & Michael B. Eisen¹

[Published in 29 Journal of Legal Studies 545 (2000)]

Abstract

Although citation to precedent in judicial opinions is a critical component of the network of rules that comprise “the law” in any area, there have been surprisingly few systematic attempts to use the abundant data available on citation patterns to uncover general principles that might illuminate the nature and structure of the legal system. In this paper, we use data from the New York Court of Appeals and the Seventh Circuit regarding the number of times judicial opinions cite to, and are subsequently cited as, precedent to test the hypothesis that legal arguments and legal doctrine have a kind of “fractal” structure. Our model provides a reasonable fit to the citation data that we examined, although there appear to be significant sources of variability in these data that are not explained by our simple predictive framework, and it is clearly far too early to draw any robust conclusions about the hypothesis other than that additional work along these lines appears to be warranted.

How Long is the Coastline of the Law?

Thoughts on the Fractal Nature of Legal Systems

David G. Post² & Michael B. Eisen³

Introduction

Citation to precedent in judicial opinions is a seriously under-studied phenomenon. It goes without saying that in a common-law system such as our own, citing precedent is one of the more significant means by which current legal disputes are resolved; indeed, one could plausibly suggest that the web of citations from one case to another is a critical component of the network of rules that comprise “the law” in any area, as any first-year law student struggling to master Shepardizing can attest. Although the concept of precedent is truly “at the heart of the way in which lawyers think about the legal system,”⁴ there have been few systematic attempts to use the abundant data available on citation patterns to uncover general principles that might illuminate the nature and structure of the legal system.⁵

In this paper, we analyze case citation patterns from a theoretical perspective derived from the theory of complex systems and fractal geometry, and we explore possible links between citation practices and a more general theory of legal decision-making and legal argumentation. More specifically, we believe that legal arguments have a kind of fractal structure — recursively generated and possessing a branching, self-similar, tree-like structure at all levels of the argumentation hierarchy — and that case citation patterns should reveal that structure. In Section I, we introduce some of the basic concepts of the geometry of fractals. Section II, building upon

a provocative early paper by Jack Balkin, introduces the notion that legal systems can be thought of as fractal objects. In Section III we formulate a simple model of a fractal-generating process, and compare the output of that model to citation data from state and federal courts (the New York Court of Appeals and the Seventh Circuit). Finally, In Section IV we discuss the results of these analyses and suggest directions for future research.

I. An Introduction to Fractal Geometry

Fractals comprise a class of geometric figures that share some rather unusual characteristics.⁶ Although there is no generally accepted definition for the term fractal, they can generally be characterized as complicated figures of infinite length that do not simplify when magnified, that is, whose structure repeats itself at all scales.⁷ To appreciate the basic nature of fractal geometry one can begin, following Benoit Mandelbrot's classic formulation,⁸ with the seemingly simple question: How do we measure the length of a geometric object?

One simple and straightforward method goes as follows: take a ruler of length ϵ , and then "walk" the ruler along the figure, starting each new ruler-step where the previous one left off. The number of ruler-steps taken ($N(\epsilon)$), multiplied by the length of the ruler (ϵ), is an estimate of the length of the figure.

$$L(\epsilon) = N(\epsilon) * \epsilon , \quad (\text{Equation 1})$$

where $L(\epsilon)$ = the estimate of length derived using a ruler of length ϵ , and

$N(\epsilon)$ = the number of steps needed to traverse the figure using a ruler of length ϵ

What happens to our estimate $L(\epsilon)$ as our measuring ruler gets smaller and smaller (that is, at higher and higher magnification)? That, as it turns out, depends critically on the kind of figure whose length we are measuring. Consider the simplest of all geometric figures, the ordinary Euclidean straight line. For a line, of course, estimated length $L(\epsilon)$ does not change as ruler

length changes (that is, $L(\epsilon) = \text{constant}$); a one-meter line remains one meter long whether $\epsilon = 1$ m (in which case $N(\epsilon) = 1$ and $L(\epsilon) = 1 * 1 \text{ m} = 1 \text{ m}$ by equation 1), 1 cm ($N(\epsilon) = 100$, $L(\epsilon) = 1 \text{ m}$), 1 mm ($N(\epsilon) = 1000$, $L(\epsilon) = 1 \text{ m}$), etc.

For an irregular figure, see Figure 1(a), the estimate $L(\epsilon)$ may change as a function of ruler length; over a certain range, as the ruler gets smaller and smaller, $L(\epsilon)$ increases, (that is, the figure appears to get longer), because smaller rulers can follow the irregular “wiggles” of the curve more closely. See Figure 1(b). But for all Euclidean figures, $L(\epsilon)$ converges to a well-defined value — the “true” length of the figure — as ruler length decreases (ϵ approaches zero).⁹

$$L(\epsilon) = N(\epsilon) * \epsilon \implies \text{constant}$$

$$\lim_{\epsilon \Rightarrow 0}$$

This is little more than a restatement of principles familiar to anyone who has taken (and remembers) high school geometry: A circle has a true length that is independent of the measuring instruments used to estimate that length; we call that length the circumference of the circle, and it is, of course, equal to the radius of the circle multiplied by π . Our estimating procedure will — must — converge on this value as our measuring rulers get smaller and smaller; we may need an infinite number of infinitesimally small rulers to get it exactly correct, but Euclidean geometry is premised on the notion that we can do just that.

True fractal objects, by contrast, are those for which (by definition) estimated length $L(\epsilon)$ never converges.¹⁰ Fractals appear to get longer and longer as the measuring stick gets smaller and smaller, and the estimated length of a true fractal diverges to infinity as ϵ approaches zero:

$$L(\epsilon) = N(\epsilon) * \epsilon \implies \infty \quad \text{[Equation 2]}$$

$$\lim_{\epsilon \Rightarrow 0}$$

A simple example is the binary tree shown in Figure 1(c), drawn by repeating a simple “forking” pattern at smaller and smaller scales. Each branch splits into two vertical branches, with a scale reduction factor of $\frac{1}{2}$. The number of vertical branches thus doubles at each level of the figure, while the length of each individual branch halves at each level. The figure created by repeating this process an infinite number of times is a “true” or “ideal” fractal.¹¹ Figure 2(d) shows a different fractal, the so-called “Koch Snowflake,” produced as shown by infinite repetition of a single geometrical transformation (as shown in Figures 2(a) - 2 (c)).

For these figures, as the measuring ruler gets smaller and smaller the estimate of the length of the object continues to increase. These figures have no “true length”; at higher and higher magnification, more of the detail — the “wiggles” — in a fractal curve appear, and no matter how small the ruler, how high the magnification, the figure continues to get longer and longer.

Fractal objects are also self-similar at all scales; that is, they have the same shape and structure at whatever scale or magnification you view them. For example, each branch of Figure 1(c), however small, can in its turn be seen as the “trunk” of a complete tree, a scaled-down copy of the entire figure — just as the tree shown in Figure 1(c) may itself represent just one branch of a much larger similarly-structured figure. If you were to “zoom in” on any portion of the figure and view it at a significantly higher level of magnification, it would look exactly the same as Figure 1(c) itself. No matter how high the magnification — no matter how “deep” into the structure you look — it always looks exactly, dizzyingly, the same.¹²

While fractal objects turn out to have a number of important characteristics of purely theoretical interest (and considerable theoretical importance),¹³ they have attracted attention in recent years primarily because many real-world objects, from semiconductors to bronchial tubes in the lung to meteorites, appear to have fractal-like qualities and to be well-described by the

geometry of fractals.¹⁴ Coastlines, for example; the irregular figure depicted in Figure 1(a) is in fact a sketch of a portion of the British coastline, but because coastlines display self-similarity at a vast range of scales you cannot tell from Figure 1(a) the scale at which it is drawn, that is, whether this drawing represents a satellite photograph (at a scale of, say, 1 cm = 1000 km), an ordinary road map (1 cm = 10 km), the outline of some particular inlet at even smaller scale (1 cm = 100 m), or the contours of individual grains of sand. Coastlines have the same “wiggly” shape at all scales, and are self-similar and scale free in that sense, and indeed, as an empirical matter, measured coastline lengths never appear to stabilize or converge on a fixed value, but keep increasing at finer and finer measurement scales. See Figure 3.

II. Is the Law a Fractal Object?

It is, perhaps, difficult to think of “the law” as a geometrical object. There is, however, a sense in which the processes of legal argumentation resemble the kind of recursive branching processes that can generate fractal curves, and in which the structure of legal doctrine possesses a fractal-like self-similarity over a wide range of scales.

Jack Balkin was, we believe, the first to make this suggestion (though it has now become unexceptional, if not quite commonplace, in the literature¹⁵). Balkin argued that legal doctrine has a branching structure similar to the fractal branching tree displayed in Figure 1(c),¹⁶ in which the “basic structure of moral and political choice is reprised at each level of discourse, so that large scale structure resembles small scale structure.”¹⁷ He illustrated this point with a well known (and quite typical) doctrinal problem from the law of torts: should a defendant’s negligence be judged against an “objective” standard (that is, what the reasonable person in defendant’s situation would have done) or a “subjective” standard (what the defendant believed was reasonable given defendant’s situation)? Any body of tort law must, in some sense, confront and resolve this

problem in some way; it represents an example of what Balkin calls a “rule choice.” Whichever choice is made, additional rule choices, at “lower” levels of the doctrinal hierarchy, will inevitably be generated. For example, under the objective standard, the question will inevitably be presented whether

“... there is an exception for children, or a different standard for insane persons, or for those who are blind, or intoxicated, and so forth. This leads us to further rule choices, each of which leads to additional branches of doctrinal development. Assume, for example, that we follow one of these branches of doctrinal development and create an exception for children (which is now the majority rule). We might consider if there is an exception to that exception when the child engages in an adult activity (this too, is the case now generally). We might then go on to ask if operating a motorcycle is an adult activity within the meaning of that rule, and if so, whether operating a motor scooter is also an adult activity.”¹⁸

And so on; this branching doctrinal structure, as Balkin notes, has no natural stopping place, but continues indefinitely downward, a “descending series of rule choices of increasing factual complexity and specificity,”¹⁹ a fractal landscape where rule choices are the branching motif, endlessly repeated at finer and finer scale. See Figure 4.

This fractal structure of legal doctrine is also reflected in (or, perhaps, itself reflects) the fractal structure of legal argumentation. Consider a typical civil lawsuit; for concreteness, assume that plaintiff is asserting a simple claim of copyright infringement against the defendant.²⁰ At the “highest” level of generality, the suit raises a single question: Is the defendant liable to the plaintiff? Substantive copyright law decomposes this single question into three (which constitute the basic “elements of the claim”): Is the work in question (call it “Work A”) protected by

copyright? Is plaintiff the owner of the copyright in Work A? Did Defendant do something to or with Work A that violates one or more of the exclusive rights of the copyright holder?

The plaintiff's complaint, in effect, is an assertion that each of these questions should be answered in the affirmative. The defendant, in her response, may move the argument "downwards" along any or all of these three branches, expanding any or all of the initial questions into a number of different sub-questions of which they are composed. For example, with regard to plaintiff's assertion that Work A is protected by copyright, defendant may expand the original unitary question into three of its sub-components and argue (a) that Work A is not an "original work of authorship" that has been "fixed in a tangible medium of expression,"²¹ (b) that Work A is uncopyrightable as a matter of law,²² and/or (c) that whatever copyright protection Work A may have had has expired. The other two original questions can be similarly expanded.

Defendant may argue that plaintiff is not the owner of the copyright in Work A because plaintiff (a) is not the "author" of Work A²³, (b) is somehow statutorily or otherwise barred from owning copyright,²⁴ and/or (c) has transferred to some third party whatever ownership rights he may have previously possessed²⁵, or defendant may argue that her actions did not violate any of the copyright holder's exclusive rights in Work A because (a) her actions are not within the scope of the copyright holder's exclusive rights, (b) she did not, in fact, take the actions plaintiff has complained of, and/or (c) her actions, though they may appear to violate one or more of the copyright holder's exclusive rights, are privileged for one reason or another.

And so on; each of these sub-questions can, in turn, support additional downward expansion.²⁶ See Figure 5. Note, too, that the argument and counter-argument may move "upwards," to higher levels in the argumentation hierarchy. Defendant may, for example, argue that the court lacks subject matter jurisdiction over the claim, that the Copyright Act does not

apply to this claim at all, and/or that the court lacks personal jurisdiction over Defendant. Each of these new branches can, in turn, be expanded “downward” in much the same way that the original branch was expanded. See Figure 6.

This branching process can, in theory, continue indefinitely.²⁷ Legal argument, in this view, like the coastline of Britain or any other fractal object, has no “true” length, no “natural” scale. Whatever the magnification scale or level of argumentation, the same kind of argumentative “wiggles” reappear. We can always “zoom in” on any argumentative point, looking at it as the “root” of a deeply branching structure, or “zoom out” and look at it as a small part of a larger recursively branching structure. The number of argumentative steps required to go from point A — the predicate facts of a case — to point B — the ultimate conclusion that Defendant is/is not liable to Plaintiff — is indeterminate a priori.

III. Testing the Hypothesis

The idea that legal doctrine and argumentation, like so much of the physical and biological world, is generated by a recursive process and has a kind of fractal structure is certainly a powerful and intriguing metaphor. Our goal here is to frame this more precisely as a testable hypothesis. What would we expect the legal world to look like if this were true?

In the real world of litigated cases, of course, the branching process does not continue indefinitely. Each decided case represents a single instantiation of this process that has come to rest at some point, each opinion a single “tree” in the forest of judicial opinions. If legal argumentation proceeds in this manner, what would that forest look like?

This requires further consideration of the workings of branching processes. We have built a simple probabilistic simulation model of a branching process that constructs simulated “trees” in

the following manner.²⁸ The model has two parameters, \underline{B} (Branch Number) and \underline{p} (Add Branch Probability), and trees are constructed in the following manner:

1. Begin with a single branch (the root);
2. Select a random number between 0 and 1. If the number selected is between 0 and \underline{p} , add \underline{B} new branches at the end-point of the branch. If the number selected is between \underline{p} and 1, terminate the branch (that is, close the end-point by placing a “leaf” at the end of the branch).
3. If the root has been terminated during Step 2, stop. If not, repeat Step 2 for all un-terminated branches (that is, for all branches that do not have “leaves” at their end-points).
4. When all branches are closed (that is, when all branches have “leaves” at all of their end-points), stop.²⁹

Fig. 7 shows one illustrative tree constructed in this manner.³⁰

How large are the trees that are produced by this simple model? How many leaves will these trees have? That, it turns out, is an interesting story. We simulated the growth of 1 million trees at each of various parameter settings, and the size distributions in these populations fall into three distinct categories. One category is displayed in Figure 8³¹, which shows the size distributions in tree populations produced by setting the parameters \underline{p} and \underline{B} so that $\underline{p} * \underline{B} < 1$, that is, so that the probability of adding \underline{B} branches at any end-point was less than $1/\underline{B}$.³² The tree-size distributions in these populations all show an orderly pattern, a steady (and approximately exponential³³) decline in the number of leaves per tree. The distributions converge rapidly to zero; the range of tree size is relatively limited (even though there is, theoretically, a finite probability that any tree in any of these three populations would have 1000, or 1,000,000, or any arbitrarily large number of leaves³⁴).

A second category of tree size distributions is produced where $p * \underline{B} > 1$, that is, where the probability of adding B branches at any end point is *greater* than $1/\underline{B}$. These distributions are not as orderly as those in the first category; instead, they appear to be bifurcated in a rather complicated way. See Figure 9. One fraction of these populations shows, as before, a steady decline in the size of the trees produced; in Figure 9(a), for example, ($p = 0.5$, $\underline{B} = 4$), the left-hand side of the distribution looks very much like the distributions displayed in Figure 8, with an approximately exponential decline in the number of trees of size 10, 50, or 100. But a second, substantial, fraction of these populations apparently never stops growing; the range of tree sizes in these populations is apparently unlimited.³⁵ In Figure 9(a), for example, approximately 45% of the trees in this population (represented by the value plotted with a ■ at $c = 1000$ in Figure 9(a)) reached a size of 10,000,000 leaves and were still growing — that is, they still had untermiated branches. The same is true for the other populations constructed with this parameter relationship, shown in Figure 9(b) - (d), each of which has a smaller, but still substantial, number of trees of (apparently) infinite size. Unlike the populations in which $p * \underline{B} < 1$, see Figure 8, the size distributions in populations in which $p * \underline{B} > 1$ do not converge to zero; it appears that once trees manage to pass some critical size in these populations, they continue to grow indefinitely.

Finally, there is a third category of distributions intermediate between these two, produced when $p * \underline{B} = 1$ — that is, where the number of branches added is the precise inverse of the probability of adding additional branches. In these populations the tree-size distribution follows a power law of the form

$$N(c) = Kc^{-x} \quad \text{[Equation 3]}$$

Where

$$K = \text{constant,}$$

c = the number of terminating end points on each tree generated by this process,

and

$N(c)$ = the number of trees with c terminating end-points.³⁶

See Figure 10.

These populations (unlike the tree populations in the first two groups) are truly “fractal” in structure. Power law distributions are characteristic of fractal objects because they embody the scale-free self-similarity that is the hallmark of those objects. They are invariant under a scale change. To be more precise, wherever we look along a power law curve the proportionality relationship of equation 3 ($N(c) \propto c^{-x}$ is undisturbed.³⁷ They are therefore scale-free; at these parameter settings the size of trees have the kind of scale-free property associated with fractal objects.

To summarize, the behavior of populations produced by the simple branching process that we simulated falls into three distinct categories or distributional domains. One category, in the region defined by $p * \underline{B} < 1$, displays a pattern of finite and orderly growth; a second, where $p * \underline{B} > 1$, a pattern of infinite, disorderly growth; and a third, at the boundary between these two regions where $p = \underline{B}$, a pattern of infinite but orderly growth at all scales in accordance with a simple power law.

This phenomenon — whereby power law behavior and fractal structure is produced at the boundary between order and disorder, at the “edge of chaos” — is a common feature of models of recursive growth processes of this kind.³⁸ Power law distributed data indicates — at least to the extent that these models accurately reflect underlying processes — that the growth pattern is poised at precisely that boundary, that somehow the parameters have been precisely “tuned” to reach that result. And while we can set the parameter settings of our computerized model at will

to produce power law size distributions (or, for that matter, distributions in either of the other two categories), one might perhaps think that distributions in this third category would rarely, if ever, be observed in the real world of complex systems; how likely is it, after all, that the parameters of a real world branching process would be so finely “tuned” as to come to rest on this narrow boundary so as to produce a power law distribution of sizes?

This, however, turns out not to be the case; power law relationships are well-nigh ubiquitous in a wide variety of physical, biological, and social systems where no “tuning” of the parameters is permitted,³⁹ leading some to suggest that there are mechanisms not yet understood whereby growth processes “self organize” and evolve without direction to this “critical” boundary state.⁴⁰

Our hypothesis, then, is that the fractal structure of systems of legal doctrine and legal argument will be reflected in a power law distribution of the output — the size, or other measurable characteristics — of those systems. Determining appropriate measures of the output of legal systems is a difficult task; as it were, before we can determine the length of the law’s coastline we need to decide precisely what the law’s coastline might be. We have chosen two measures of the “size” of the output of the legal system: (a) the number of citations in judicial opinions to previously-decided cases, and (b) the number of citations *to* those judicial opinions in subsequently-decided cases.

To some extent, of course, our decision was driven by practical considerations; case citation data is widely available in various electronic databases in a form that permits relatively easy analysis of large data-sets. But we believe that the number of citations to prior precedent contained in a judicial opinion, and the number of subsequent citations to an opinion, are important (though they are surely not the only) components of the output of our legal system and

help define the structure of that system. The number of citations to prior precedent in any judicial opinion should provide at least some information about, and can serve as a rough proxy for, the “size” of legal arguments (at least, the size of those arguments as they are ultimately formulated by the judge(s) deciding the case); while citations serve many functions, their primary function is surely to point to a prior case as having resolved a particular legal issue raised by the case then under consideration in a particular way — to provide the answer to a question posed by the current controversy. The difference between a case that cites to 100 previously decided cases and one that cites to 10 is surely due to many factors — “the judge’s personal style, tastes, erudition, pedantry, etc.”⁴¹ — but we think it reasonable to assert that the former case in some sense raises (and resolves) more questions than the latter.

And the number of times that individual opinions do, or do not, show up as citations in subsequent cases should provide at least some information about the structure of the overall doctrinal web and the significance of individual cases within that web. Cases cited with great frequency are in some sense more firmly embedded within the legal firmament, having spawned more “offspring,” than those less-frequently cited that die leaving few descendants.

We chose, as our population of reported cases, all cases decided in 1930, 1950, 1970, and 1980 by the New York Court of Appeals (Cohorts 1 - 4), and all cases decided in those same years by the federal Court of Appeals for the Seventh Circuit (Cohorts 5 - 8).⁴² Our results are displayed in Figures 11 - 14.⁴³

These results provide some support for our hypothesis. On the one hand, the power law function provides a statistically significant fit to the observed data in all cohorts (and across cohorts); that is, the hypothesis that these data are distributed in accordance with a power law explains a statistically significant amount of the variation in citation practices in all cases. On the

other hand, it is apparent that additional factors are at work in producing the observed patterns. First, there appear to be clear and systematic deviations from the log-linear form predicted by the power law function in many of the populations, an occasionally pronounced curvature more suggestive of an exponential, rather than a power law, distribution.⁴⁴ This is particularly noticeable in the federal court citation data, both for the distribution of the number of citations in opinions, see Figure 16, and the number of times opinions are subsequently cited, see Figure 18. In all of these 7th Circuit populations, there appears to be an almost constant probability of “small” events (that is, cases with relatively low numbers of citations, or opinions cited few times); these distributions appear almost horizontal in the range $c < 10$ or so, with the predicted log-linear trend showing up only in the range of $c > 10$. The New York data, on the other hand, much more closely approximate the expected linear shape across the board — and to the extent they deviate from that prediction, they appear to do so in a different way (that is, there appear to be *more* small events than predicted by the power law model).

IV. Discussion

We regard the results of this preliminary study as highly suggestive but hardly conclusive. The fractal structure hypothesis accounts for a limited (though statistically significant) fraction of variation in citation practices in the data that we examined — enough, we believe, to warrant further research along these lines.⁴⁵ There are any number of possible directions such research might take. We can certainly imagine more sophisticated tests of the hypothesis. There may well be far better measures of legal system output than the case citation statistics we chose to focus on here. Moreover, our analysis can hardly be considered definitive even as applied to those citation data. We opted for quantity rather than quality; we did no “clean up” of the raw citation data, lumping together all citation types (“See,” “But see,” “See also,” “Cf.,” etc.) into a single

measure of number of citations per opinion, nor did we attempt to eliminate, or correct for, the practice of string citations. Used in this fashion, the raw number of citations per case is undoubtedly an extremely crude proxy for the variable of interest, the number of different legal questions resolved by the opinion, and it is likely that a more fine-grained and careful analysis of these data would yield a better understanding of the extent to which this hypothesis meaningfully accounts for observed patterns in these systems.

More generally, we view these analyses as but the initial, halting steps in what we believe is a larger and more significant overall project. Some forty-odd years ago, Lon Fuller suggested that we think about the legal system as an “imperfectly growing tree” that, with pruning, can “realize its own capacity for perfection [and] grow properly.”⁴⁶ Fuller can be forgiven for not considering the possibility that legal systems actually — and not just metaphorically — grow as trees do, for it is only in recent years, with the development of fractal geometry and the other quantitative tools for studying complex systems, that we are beginning to understand the general principles that may govern the growth process in trees and other natural objects.⁴⁷ We do not, as we have said, regard this question as having been in any sense answered in this paper; we do believe that our results suggest that it is worth continuing to ask it. If, indeed, there are principles common to the processes of growth in legal systems and in other biological and physical systems, the implications for the study of the law would seem to be substantial⁴⁸; like other powerful paradigms (law and economics comes to mind), this view of the nature of the legal universe generates a series of questions — can we identify particular configurations of legal systems that are more, or less, stable? How does the size and complexity of legal systems change through time? Are there systematic differences between different bodies of law in terms of their “fractal

dimensions”? How does the law “colonize” new areas? — derived from this general theory that can perhaps be answered with (and only with) these tools.

Table 1. Number of Reported Opinions in Each Cohort

Cohort	Description	No. of opinions
1	NY Court of Appeals, 1930 (vols 252-255 NY)	918
2	NY Court of Appeals, 1950 (vols 300 - 302 NY)	1313
3	NY Court of Appeals, 1970 (vols 25 - 27 NY2d)	733
4	NY Court of Appeals, 1980 (vols 48 - 52 NY2d)	1281
5	7 th Circuit, 1930 (vols 37 - 45 F2d)	181
6	7 th Circuit, 1950 (vols 179 - 186 F2d)	213
7	7 th Circuit, 1970 (vols 421 - 440, F.2d)	434
8	7 th Circuit, 1980 (vols 615 - 645 F.2d)	410

Note that the correspondence between the year of the decision and the data in each cohort is only approximate. Because of the way in which the data in the Shepard's database are stored, it was considerably easier to extract information based upon the volume number of the reporter in which the opinion was published, rather than the year of the decision. We constructed each cohort by taking the first and last volumes containing any cases decided in a particular year, and examining all cases in those volumes. For example, there are no N.Y. Court of Appeals cases

with a decision date between January 1 and December 31, 1930, in any volume of N.Y. Reports prior to volume 252, or subsequent to volume 255; accordingly, Cohort 1 consists of all cases in volumes 252-255.

Figure legends.

Figure 1. (a) An irregular curve. (b) Estimating the length of the irregular curve shown in Figure 1(a), using a ruler of length E . (c) A fractal branching tree, created by recursive addition of a two-branch “fork” at all branch end-points.

Figure 2. Constructing the “Koch Snowflake.” Beginning with a line of unit length (a), place an equilateral triangle over the middle third of the line (b). Repeat for each of the four line segments in the new figure (c), and so on (d).

Figure 3. Empirical data on the length of various coastlines and land boundaries. These data are derived from Mandelbrot, *supra* note 8, who in turn derived them from L. F. Richardson, The Problem of Contiguity: An Appendix to *Statistics of Deadly Quarrels*, 6 Gen. Sys. Yearbook 139 (1961); other coastline data, and a discussion of these “Richardson plots,” can be found in Lauwerier, *supra* note 6, at 29-31, and Williams, *supra* note 6, at 309-13. Note that the length of the circle plotted in this figure converges rapidly to a constant value.

Figure 4. A branching diagram of the “rule choice” between different forms of tort doctrine. From Balkin, Crystalline Structure, *supra* note 16.

Figure 5. The elements of the copyright claim (P = plaintiff, D = defendant). See text for details.

Figure 6. The “copyright tree” shown in Figure 5 displayed within a larger superstructure. See text for details.

Figure 7. An example of a tree constructed by our simulation model, containing 10 “leaves” (terminating branches, represented by closed circles). At each rectangle, the program flips a (virtual) coin; if the coin comes up heads — which it does with probability p — the program adds B branches (here, two). If the coin comes up tails, the program terminates the branch.

Figure 8. Each graph in Figure 8 shows the size distribution in a population of simulated trees at the stated parameter settings. The x-axis represents the number of “leaves,” c , and the y-axis the number of trees in the population having exactly c leaves ($N(c)$). In all populations shown in this figure, the model parameters were set so that $\underline{p} * \underline{B} < 1$. On each graph we have plotted the least-squares regression line for an exponential distribution (along with the r-squared value for the regression).

Figure 9. Graphs of the distribution of tree sizes in populations with $\underline{p} * \underline{B} > 1$. Graphing conventions as in Figure 8. The solid square plotted at the value of $c = 1000$ represents the number of trees that were still growing when we terminated their growth at 10,000,000 leaves.

Figure 10. Graphs of the distribution of tree sizes in populations with $\underline{p} * \underline{B} = 1$. Graphing conventions as in Figure 8. On each graph we have plotted the least-squares regression line for a power law distribution (along with the r-squared value for the regression).

Figure 11. The distribution of the number of citations within individual cases within our four cohorts ((a) — (d)) from the New York Court of Appeals, and for the aggregate New York data from all four cohorts (e). On each graph we have plotted the least-squares regression line for a power law distribution (along with the r-squared value for the regression).

Figure 12. The distribution of the number of citations within individual cases within our four cohorts ((a) — (d)) from the Seventh Circuit Court of Appeals, and for the aggregate Seventh Circuit data from all four cohorts (e). On each graph we have plotted the least-squares regression line for a power law distribution (along with the r-squared value for the regression).

Figure 13. The distribution of the number of times that opinions are cited as precedent by subsequently-decided cases, within our four cohorts ((a) — (d)) from the New York Court of Appeals, and for the aggregate New York data from all four cohorts (e). On each graph we have

plotted the least-squares regression line for a power law distribution (along with the r-squared value for the regression).

Figure 14. The distribution of the number of times that opinions are cited as precedent by subsequently-decided cases, within our four cohorts ((a) — (d)) from the Seventh Circuit Court of Appeals, and for the aggregate Seventh Circuit data from all four cohorts (e). On each graph we have plotted the least-squares regression line for a power law distribution (along with the r-squared value for the regression).

Figure 1(a)




Figure 1(b)




Figure 1(c)




Figure 2

(a)



[scale = 1 unit]

(b)




[scale = 1 unit]

(c)



[scale = 1 unit]

(d)



[scale = 1 unit]

Figure 3

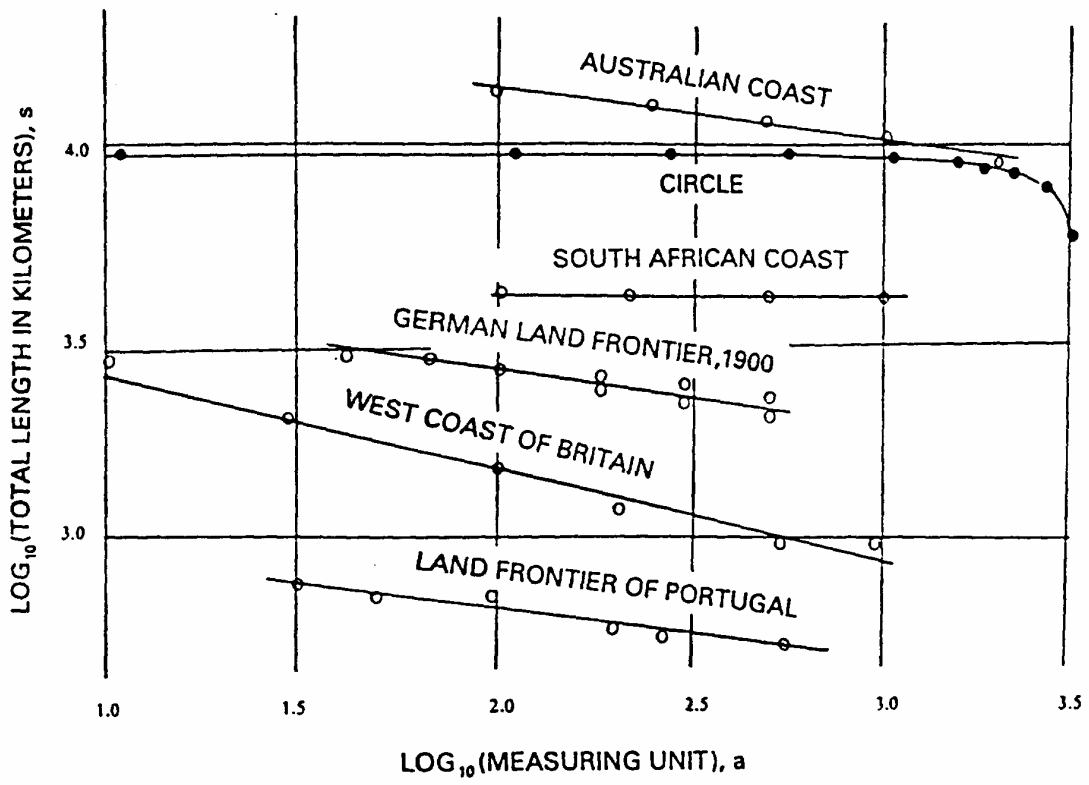


Figure 4

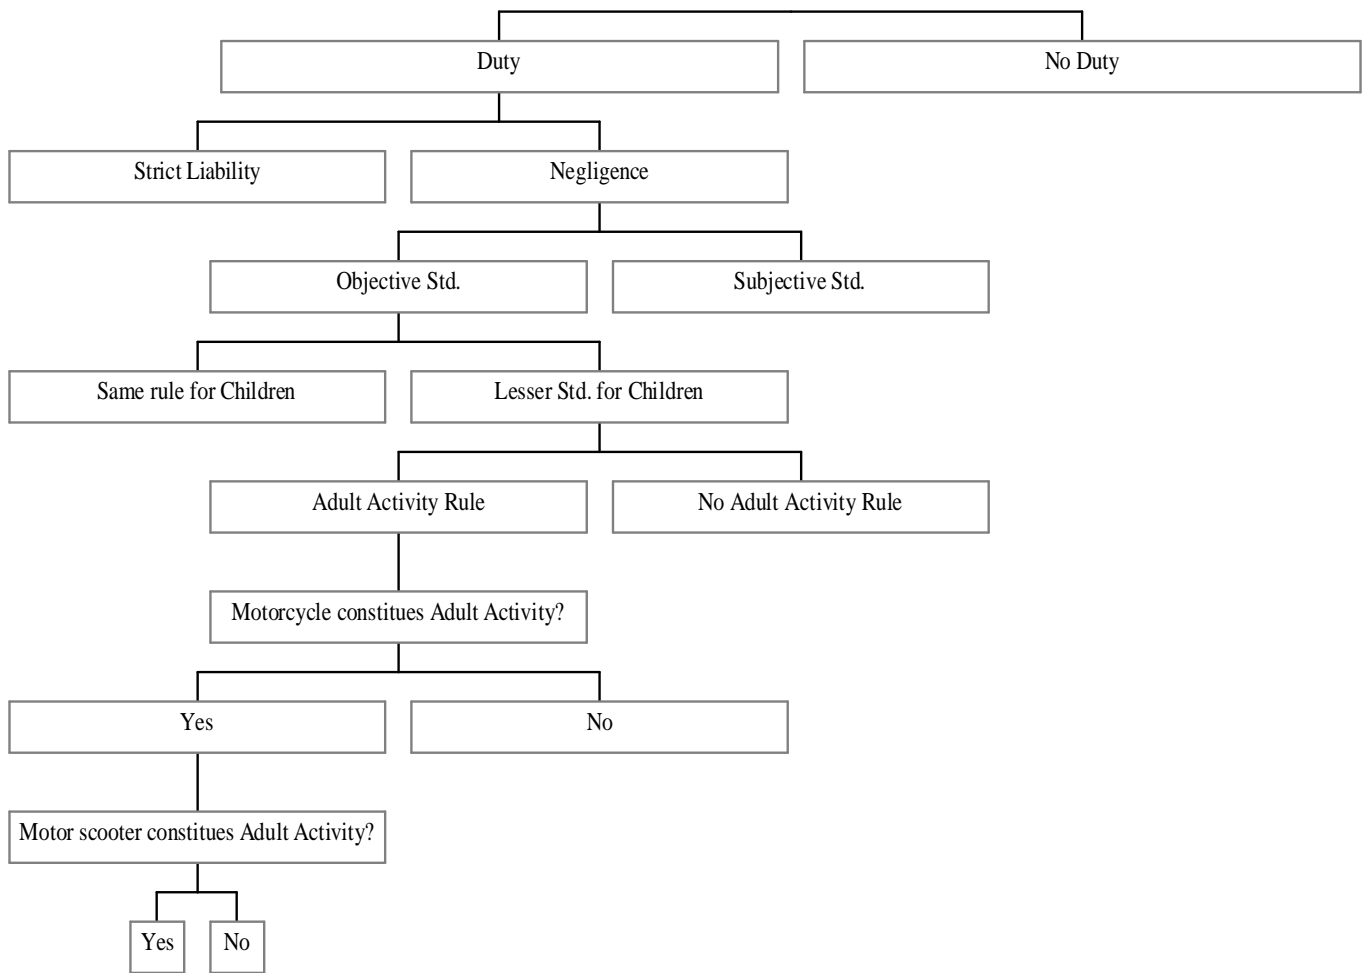


Figure 5.

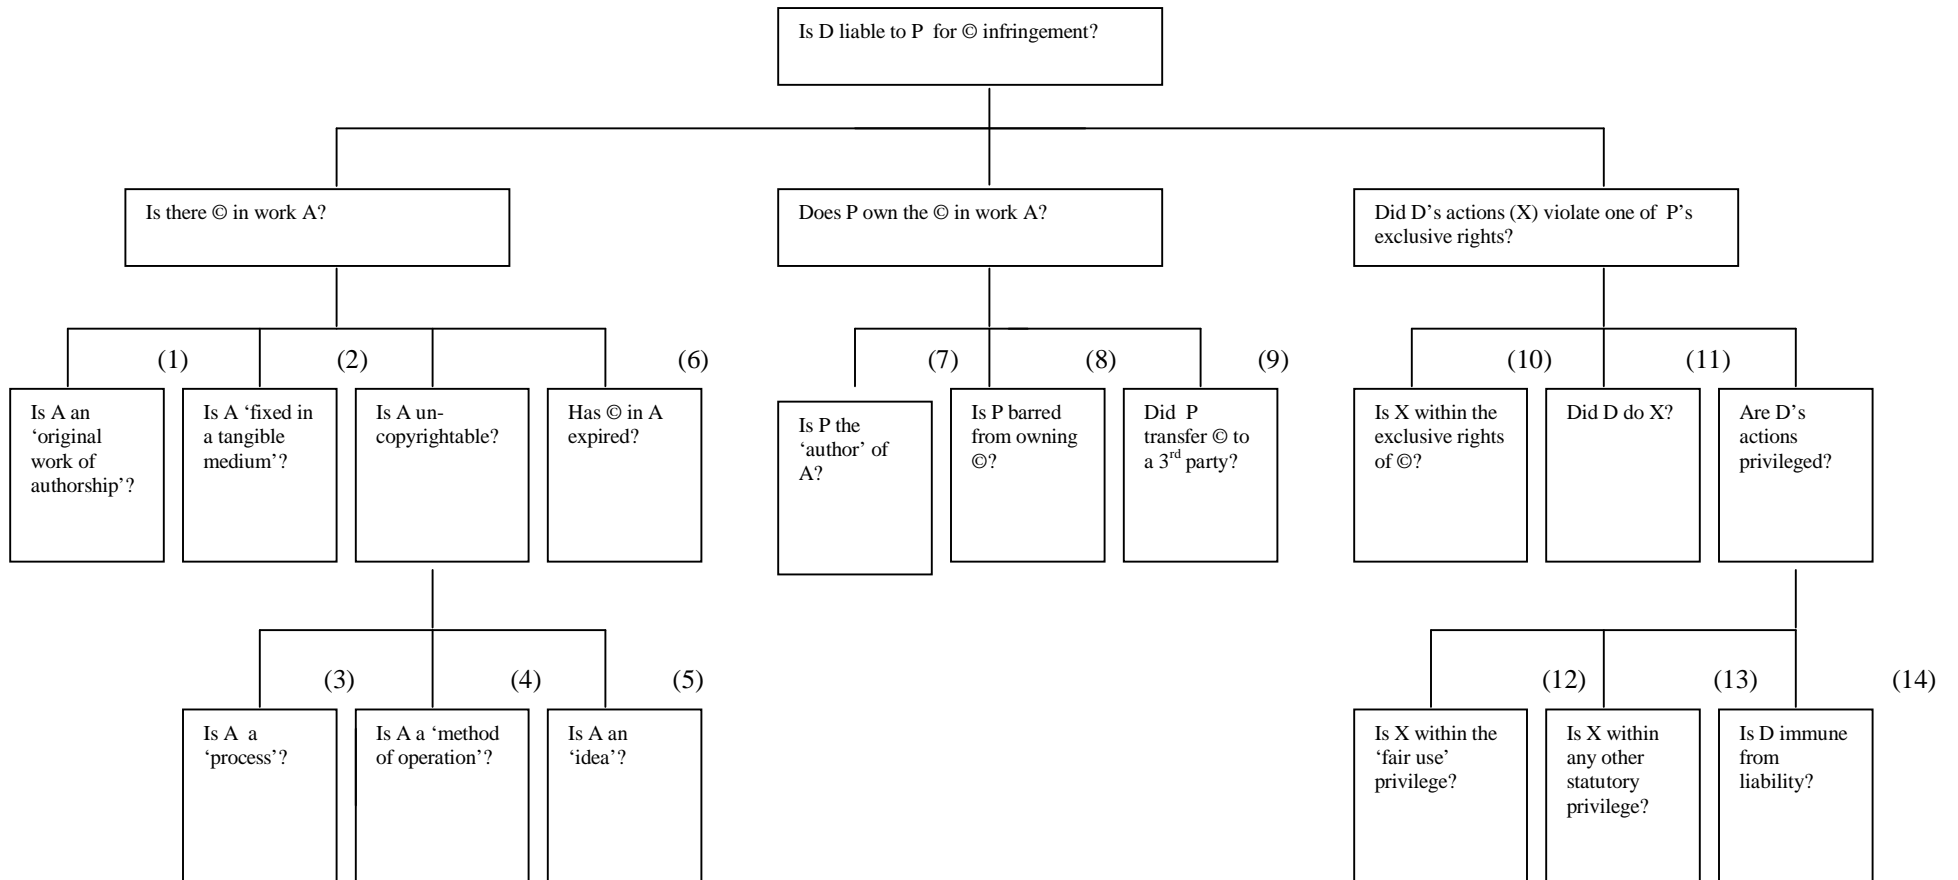


Figure 6




Figure 7




Figure 8(a).




Figure 8(b).


Branching simulation, $P=0.25$, $B=2$ 

Figure 8(c).




Figure 9(a)




Figure 9(b)




Figure 9(c)




Figure 9(d)




Figure 10(a)




Figure 10(b)

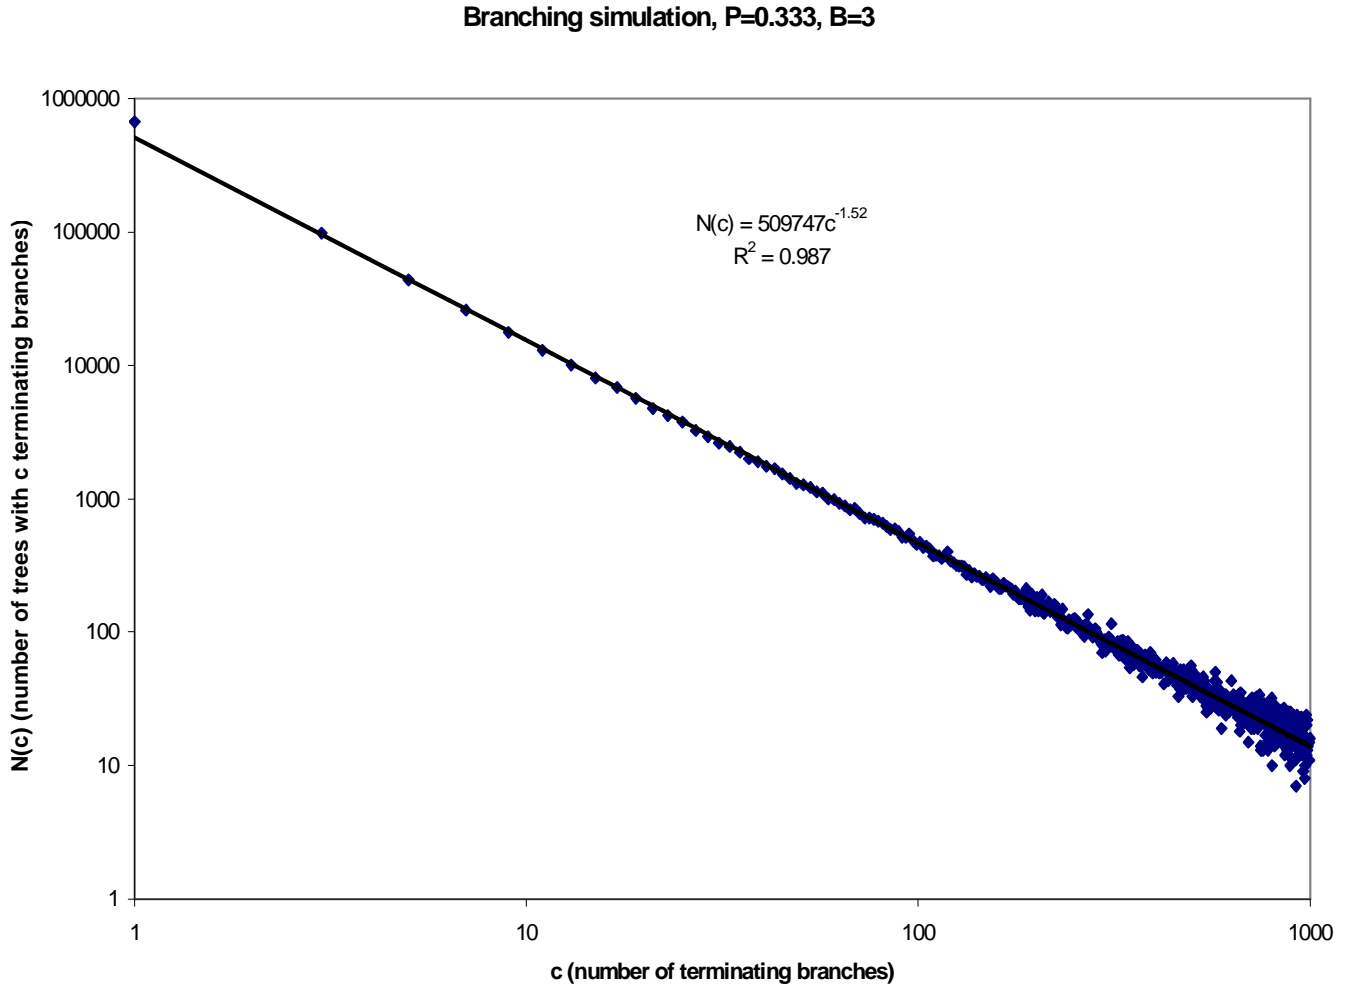


Figure 10(c)




Figure 10(d)

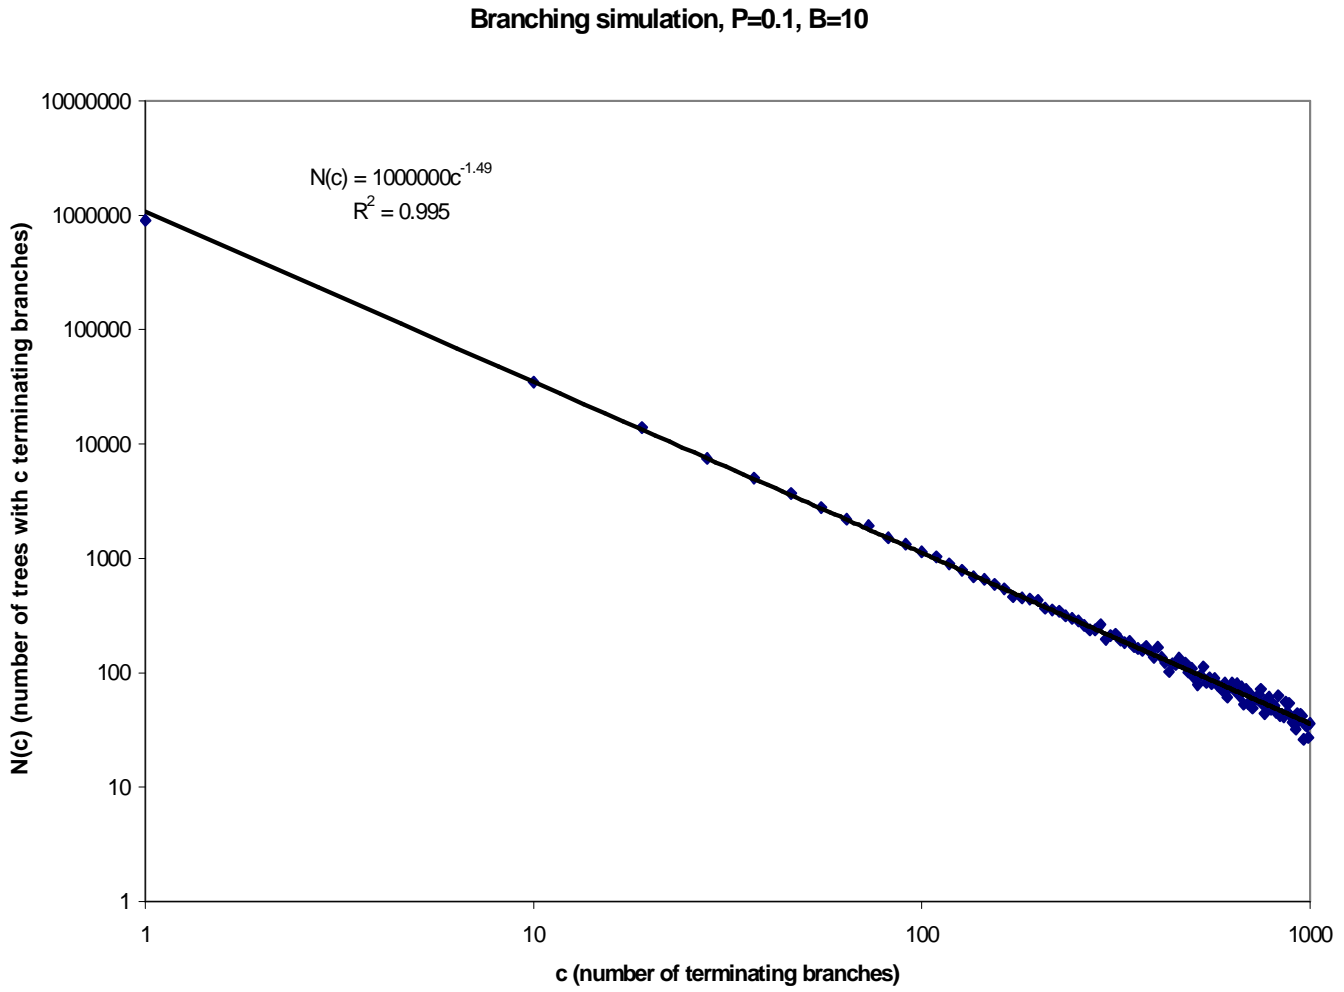


Figure 11(a)




Figure 11(b)




Figure 11(c)




Figure 11(d)

Cohort 4 (NY Ct. App. 1980)




Figure 11(e)

All NY Ct. App. Data




Figure 12(a)




Figure 12(b)




Figure 12(c)

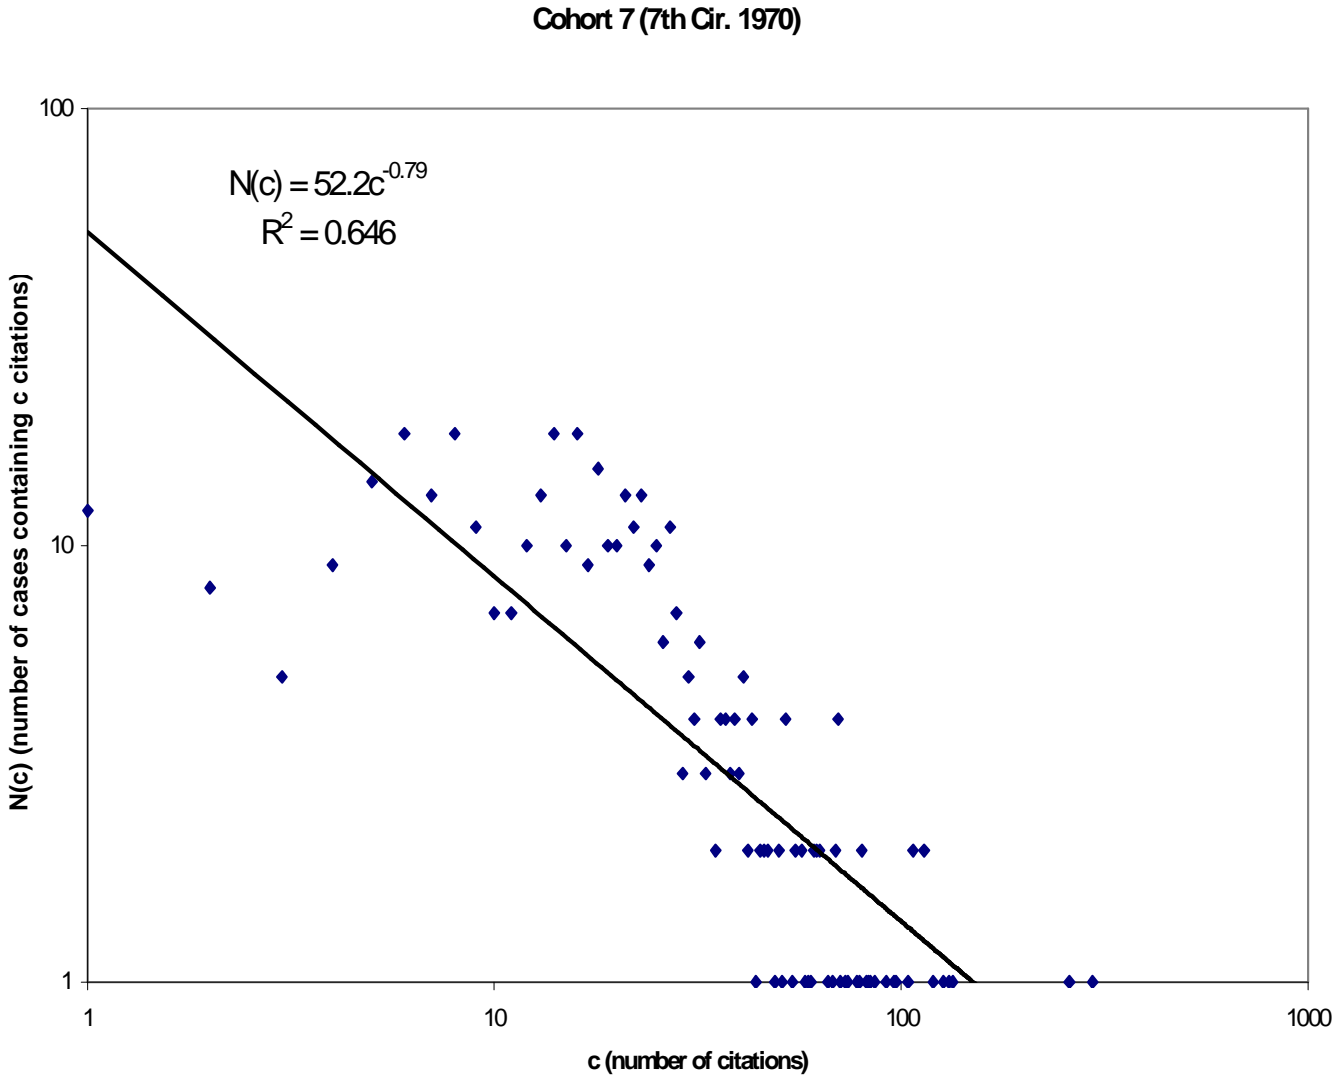


Figure 12(d)




Figure 12(e)

All 7th Cir. Data

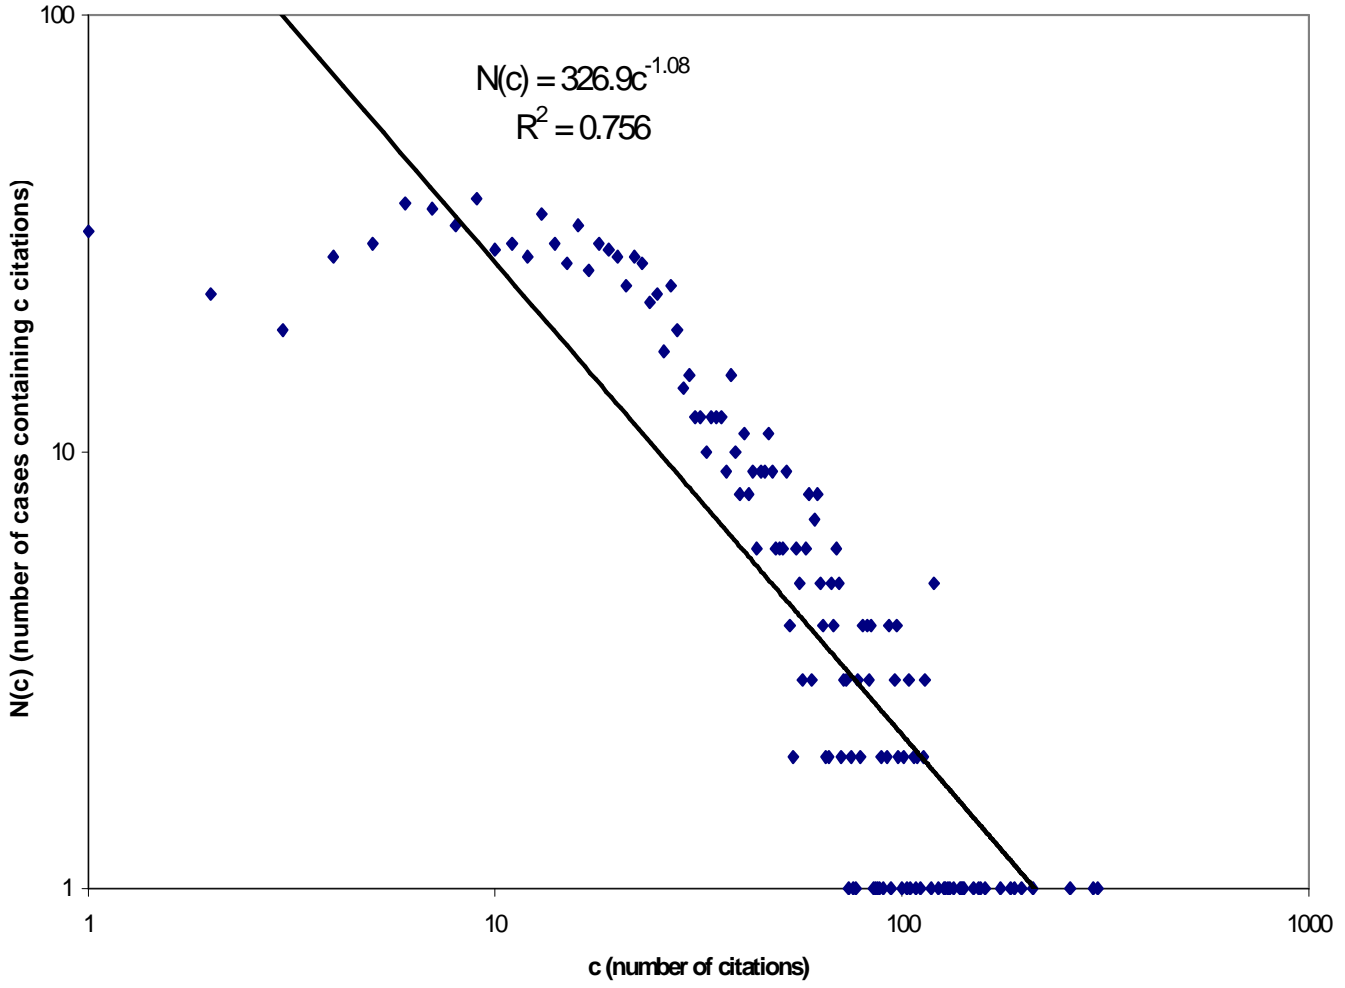


Figure 13(a)




Figure 13(b)

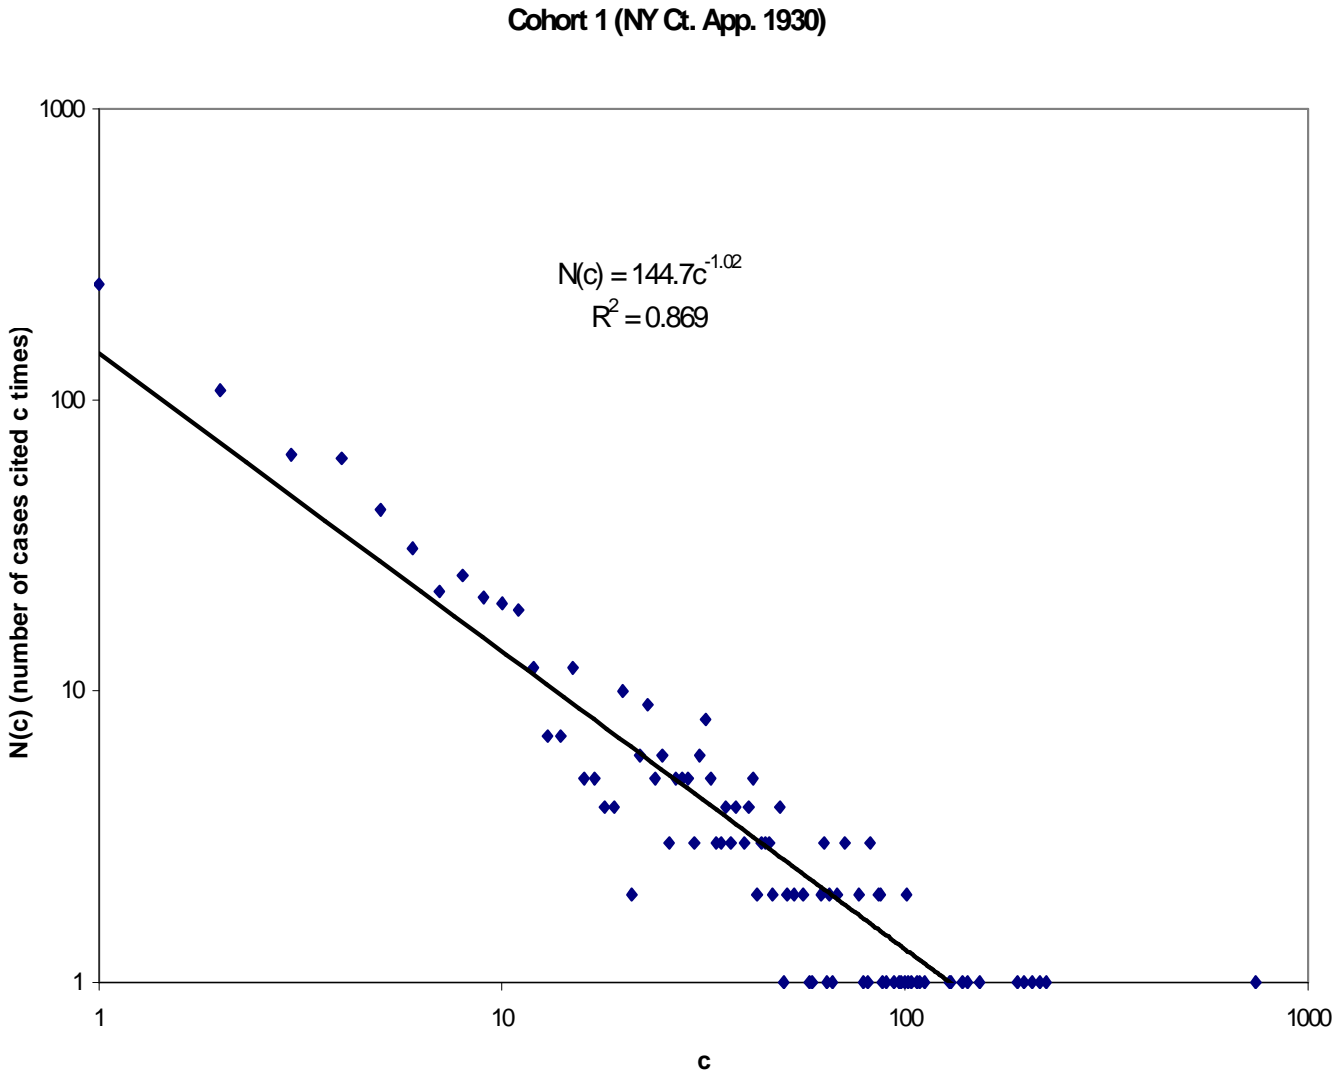


Figure 13(c)

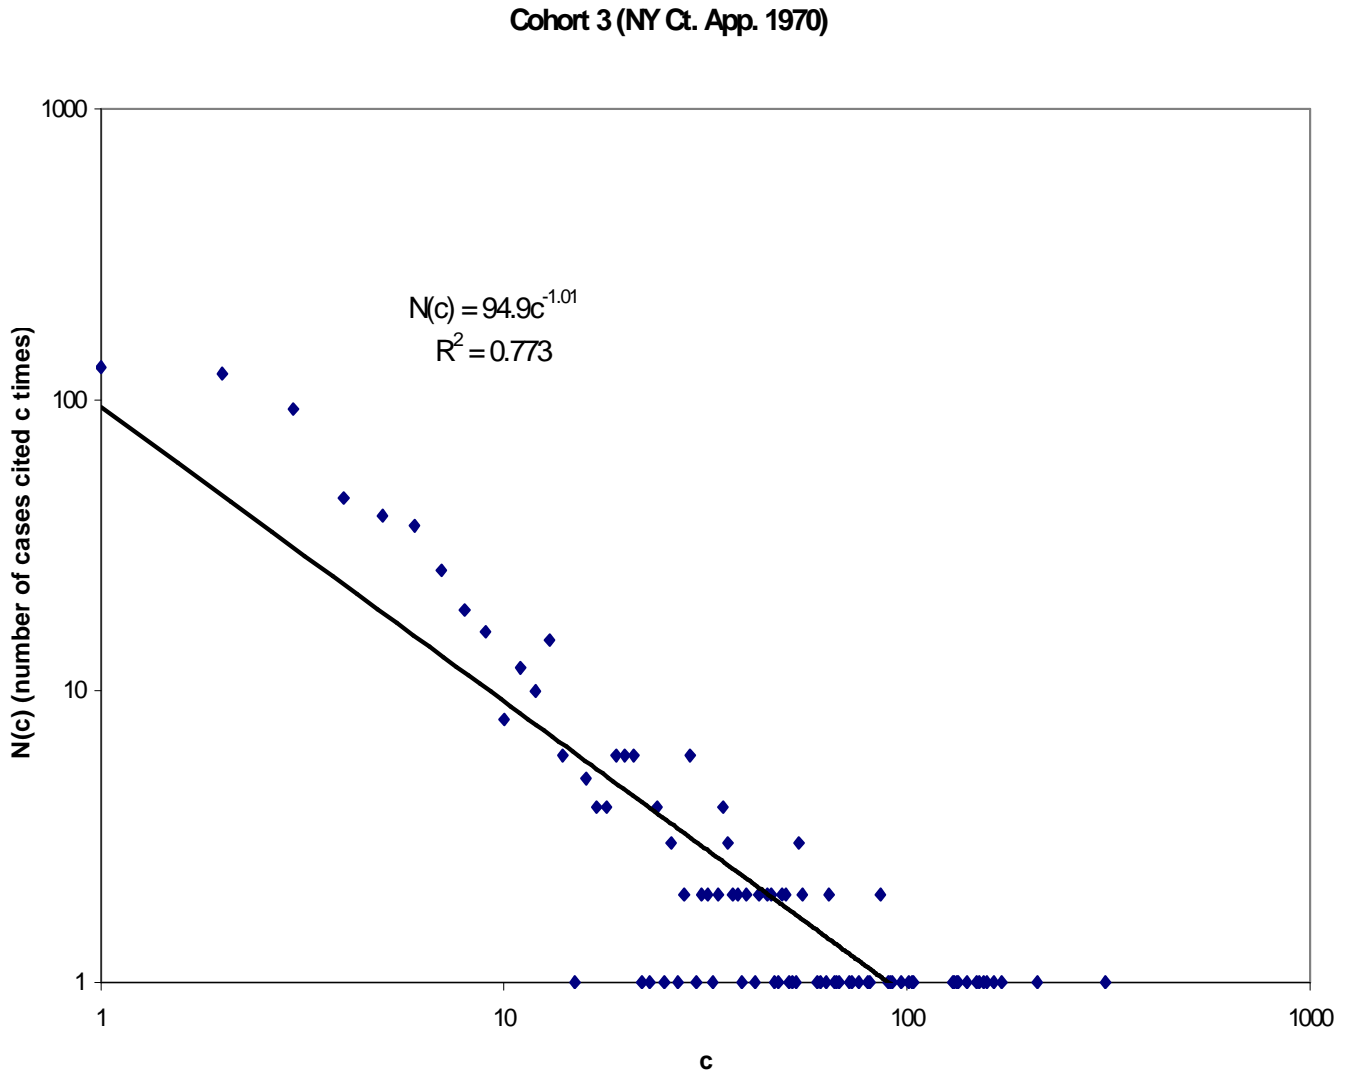


Figure 13(d)

Cohort 4 (NY Ct. App. 1980)




Figure 13(e)




Figure 14(a)




Figure 14(b)




Figure 14(c)




Figure 14(d)





Figure 14(e)



ENDNOTES

¹We gratefully acknowledge the assistance provided by Gary Spivey, Mark Briody, and their colleagues at Shepard's, Inc., in providing us with the citation data reported on in this paper. We

also thank David Johnson and Steve Salop for many enlightening discussions about the fractal nature of legal argumentation, Dawn C. Nunziato for a very helpful reading of an earlier draft, the participants at the Northwestern University Conference on Legal Citation for their comments on our presentation of this paper, and Larry Lessig and Fred Shapiro for arranging for our participation at that Conference and in this symposium.

²Beasley School of Law at Temple University; Dpost@vm.temple.edu.

³Department of Genetics, Stanford University Medical School.

⁴William M. Landes and Richard Posner, "Legal Precedent: A Theoretical and Empirical Analysis," 19 J. Law & Econ. 249, 292 (1976).

⁵There are, of course, exceptions to that general rule. See, for example, William M. Landes and Richard Posner, *supra* note 4; William M. Landes and Richard Posner, "Legal Change, Judicial Behavior, and the Diversity Jurisdiction," 9 J. Legal Studies 367 (1980); Michael S. Fried, The Evolution of Legal Concepts: The Memetic Perspective, 39 Jurimetrics 291 (1999); William M. Landes, Lawrence Lessig, and Michael E. Solimine, "Judicial Influence: A Citation Analysis of Federal Courts of Appeals Judges," 27 J. Legal Studies 271 (1998); Virginia Cano, "Citation Life Cycles of Ten Citation Classics," 22 Scientometrics 297 (1991); John Merryman, "Toward a Theory of Citations: An Empirical Study of the Citation Practice of the California Supreme Court in 1950, 1960, and 1970," 50 S. C. L. Rev. 381 (1977); Peter Clinch, "The Use of Authority: Citation Patterns in the English Courts," 46 J. Documentation 387 (1990); James Leonard, "An Analysis of Citations to Authority in Ohio Appellate Decisions Published in 1990," 1994 Law Library Journal 129; Stephen Marx, "Citation Networks in the Law," 10 Jurimetrics 121 (1970); Fred Shapiro, "Origins of Bibliometrics, Citation Indexing, and Citation Analysis: The Neglected

Legal Literature,” 43 J. Am. Soc. Info. Sci. 337 (1992). Nonetheless, we think it fair to suggest that there has been surprisingly little work given the critical position of citation to precedent within the Anglo-American legal system.

⁶There is, to put it mildly, an immense literature about fractal geometry. Fortunately, there are a number of excellent non-technical introductions to the field: in roughly ascending order of technical difficulty, see, for example, Garnett P. Williams, *Chaos Theory Tamed* (1997); G. Sugihara and R.M. May, *Applications of Fractals in Ecology*, 5 *Trends in Ecol. Evol.* 79 (1990); Hans Lauwerier, *Fractals: Endlessly Repeated Geometrical Figures* (1991); Gary William Flake, *The Computational Beauty of Nature: Computer Explorations of Fractals, Chaos, Complex Systems, and Adaptation* (1998); Benoit Mandelbrot, *The Fractal Geometry of Nature* (1983); Manfred Schroeder, *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise* (1991); Kathleen T. Alligood, Tim D. Sauer, and James A. Yorke, *Chaos: An Introduction to Dynamical Systems* (1996); Robert L. Devaney, *Chaos, Fractals, and Dynamics* (1990). The material in this Section was compiled largely from these references. There is also an excellent collection of Frequently Asked Questions about fractals compiled by the moderators of the sci.fractals newsgroup at <ftp://rtfm.mit.edu/pub/usenet/news.answers/fractal-faq>.

⁷On the absence of a generally accepted definition for the term “fractal,” see Alligood *et al.*, *supra* note 6, at 149 - 50 (“Scientists know a fractal when they see one, but there is no universally accepted definition [although it] is generally acknowledged that fractals have some or all of the following properties: complicated structure at a wide range of length scales, repetition of structures at different length scales (self-similarity), and a ‘fractal dimension’ that is not an integer.” [See note 10 *infra* for a discussion of “fractal dimension”]. See also Williams, *supra*

note 6, at 237 - 243. Benoit Mandelbrot, who coined the term “fractal” in a 1975 essay, has suggested that we would be better off without a precise definition of the term. See Mandelbrot, *supra* note 6, at 361 - 363.

⁸Benoit Mandelbrot, “How long is the coast of Britain? Statistical self-similarity and fractional dimension,” 155 *Science* 636 (1967).

⁹See Mandelbrot, *supra* note 6, at 25 - 33; Schroeder, *supra* note 6, at 9 - 11.

¹⁰*Id.* More generally, Equation 2 states that for fractals $L(\epsilon) = N(\epsilon) * \epsilon^D$ diverges to infinity if $D = 1$. One of the defining characteristics of true fractals is that there exists a critical exponent ($D = D_{crit} > 1$), for which the expression

$$L(\epsilon) = N(\epsilon) * \epsilon^D$$

does converge to a finite number. D_{crit} is the so-called “Hausdorff,” or fractal, dimension of the object concerned. For such objects, the expression $L(\epsilon) = N(\epsilon) * \epsilon^D$ diverges to infinity for all $D < D_{crit}$, and converges on zero for all $D > D_{crit}$. See Schroeder, *supra* note 6, at 9 - 16, 213-14; Williams, *supra* note 6, at 306 - 333.

¹¹Williams, *supra* note 6, at 242 - 43, usefully distinguishes between what he calls “deterministic” or “exact” fractals (such as the fractals depicted in Figures 3 and 4), formed by (theoretically) infinite repetition of an equation or geometrical motif, and “natural” fractals, which contain elements of randomness or noise and which are necessarily limited to a finite range of sizes. See also Schroeder, *supra* note 6, at 30.

¹²Fractal objects are thus entirely *internally scale-free*; there is no way to determine the scale of the object without reference to something external to the fractal itself, precisely because it looks exactly the same at all scales. Fractal objects therefore can have no “true” or “natural” scale, no

degree of magnification that is “better” than any other for showing the shape and details of the object.

¹³Fractals have been characterized as “mathematical monsters,” Schroeder, *supra* note 6, at 8, because of some of their rather peculiar properties. They are continuous curves that, while confined within a finite area — the edge of the branching tree shown in Figure 3 will never reach the end of the page, just as a Koch Snowflake beginning with a line of unit length will never go beyond the original end-points of the unit line — are of infinite length, a seeming contradiction. They are also nowhere differentiable along that infinite length (that is they are tangent-less).

Mandelbrot describes the

“ . . . conventional attitude towards [the Koch Snowflake] on the part of mathematicians. They are all but unanimous in proclaiming that [it] is a monstrous curve! . . . [The eminent mathematician] Charles Hermite [wrote] of ‘turning away in fear and horror from this lamentable plague of functions with no derivatives.’ . . . Not only near every book but every science museum proclaims that nondifferentiable curves are counter-intuitive, ‘monstrous,’ ‘pathological,’ or even ‘psychopathic.’”

Mandelbrot, *supra* note 6, at 33. These “lamentable functions,”

. . . continuous but without tangents, were first defined a century ago by the German mathematician Karl Weierstrass, just to show his skeptical colleagues . . . that such functions did indeed exist. But other authorities, not least the great Austrian physicist Ludwig Boltzmann, saw the light: Boltzmann wrote to Felix Klein (in 1898) that nondifferentiable functions could have been invented by physicists because there are problems in statistical mechanics ‘that absolutely necessitate the use of nondifferentiable

functions.’ And his French colleague Jean Perrin went even further when, in 1906, he presaged present sentiment about such mathematical monsters, saying that ‘curves that have no tangents are the rule, and regular curves, such as the circle, are interesting but quite special.’”

Schroeder, *supra* note 6, at 7 - 8.

¹⁴ See, generally, the sources cited *supra*, note 6.

¹⁵ See, for example, Robert Scott, Chaos Theory & The Justice Paradox, 35 Wm. & Mary L. Rev. 329, 348-49 (1993) (suggesting that the conflict between future justice and present justice produces an inevitable tension in the law that replicates itself over and over, producing a fractal structure); Llewellyn J. Gibbons, No Regulation, Government Regulation, or Self-Regulation: Social Enforcement or Social Contracting for Governance in Cyberspace, 6 Cornell J. Law & Pub. Pol’y 475 (1997) (discussing the fractal nature of Internet communication and Internet rule-making). A search of the LEXIS Lawrev database revealed 55 references to the law’s “fractal structure” in published articles.

¹⁶ Jack Balkin, The Crystalline Structure of Legal Thought, 39 Rutgers L. Rev. 1 (1986) (hereinafter Balkin, *Crystalline Structure*). In a later article — Jack Balkin, The Promise of Legal Semiotics, 69 Tex. L. Rev. 1831, 1835-6 (1991) — Balkin himself noted that the argumentation structure he outlined in 1986 was better described as “fractal” rather than “crystalline.”

¹⁷ Jack Balkin, The Promise of Legal Semiotics, 69 Tex. L. Rev. 1831, 1836 (1991).

¹⁸ Balkin, *Crystalline Structure*, *supra* note 16, at 11.

¹⁹ Balkin, *Crystalline Structure*, *supra* note 16, at 9-10.

²⁰ Our choice of a copyright lawsuit for this illustration reflects only our greater familiarity with that body of law; this exercise can, we would assert, be replicated in any body of legal doctrine.

²¹ Section 102(a) of the Copyright Act, 17 U.S.C. §102(a), provides that “[c]opyright protection subsists . . . in original works of authorship fixed in any tangible medium of expression.”

²² Section 102(b) of the Copyright Act, 17 U.S.C. §102(b), qualifies section 102(a), see *supra* note 21, by stating that “[i]n no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation concept, principle, or discovery . . .”

²³ Section 201(a) of the Copyright Act, 17 U.S.C. §201(a), provides that “ownership of a copyright vests initially with the author.”

²⁴ See, e.g., section 105 of the Copyright Act, 17 U.S.C. §105, which provides that “[c]opyright protection under this title is not available for any work of the United States Government.”

²⁵ A series of rules regarding the transfer of copyright ownership is provided in section 204 of the Copyright Act, 17 U.S.C. §204.

²⁶ For example, the assertion that Work A is uncopyrightable as a matter of law can be expanded to an argument that Work A is an ‘idea’ or a ‘process’ or a ‘method of operation,’ each of which is deemed uncopyrightable in section 102(b) of the Copyright Act, 17 U.S.C. §102(b). See Figure 5.

²⁷ See Balkin, *supra* note 16, at 12 (suggesting that there can be “an infinite number of subdoctrinal choices which could follow ‘beneath’ any rule choice”). Although we believe that the assertion is correct, it may not be provable in any rigorous way; the most we can say, perhaps, is that we have never met a legal question that could not be decomposed into sub-questions, if only

because the inherent ambiguity of language is such that we can always ask, with respect to any question, what the individual words mean (and what the words comprising their definitions mean, and so on).

Prof. Landes's model of issue resolution during the litigation process is similar to the branching process we describe above. See William M. Landes, *Sequential Versus Unitary Trials: An Economic Analysis*, 22 J. Leg. Stud. 99, 124 - 132 (1993). Landes asks how the identification of separate issues in a case, and the order in which those issues are decided by the court, affects the costs of litigation and the parties' litigation strategies. He notes at the outset that "since the boundaries of an issue [in any case] are sometimes arbitrary, one may be able to separate one issue into several narrower issues," *id.* At 124, and indeed that "one may be able to subdivide issues ad infinitum." *Id.*, at 124 n. 46. He suggests, however, that restricting the definition of an "issue" to a question that is dispositive of the entire case "may limit this [infinitely recursive] process." *Id.*

Although we have made a similar suggestion ourselves, see David G. Post and Steven C. Salop, *Issues and Outcomes, Guidance, and Indeterminacy: A Reply to Professor John Rogers and Others*, 49 Vand. L. Rev. 1069, 1077 – 1081 (1996) (defining, for purposes of determining the optimal appellate court voting procedure, a "primary issue" as one that (a) is logically independent of other issues presented, (b) is potentially dispositive of the outcome of the case, and (c) cannot be decomposed further into separate sub-questions that meet criteria (a) and (b)), we are no longer persuaded that this restriction limits this infinitely branching process. For instance, the "root" question in our hypothetical copyright case – is defendant liable to plaintiff? – is a dispositive issue; if plaintiff prevails on that issue, he wins. The issues into which it can be

decomposed are also dispositive issues; if defendant can show that copyright does not protect the work in question, she prevails, or that plaintiff does not own the copyright, she wins. The issues into which those sub-questions can be decomposed are also dispositive issues; if defendant can show that the work is not an “original work of authorship,” or that copyright in the work has expired, she wins. The issues into which those sub-sub-questions can be decomposed, see Figure 5, are also dispositive issues; if defendant can show that the work in question is a “process,” or that her actions are within the fair use privilege, she wins. Restricting attention to only those issues that can be dispositive in any case does not, it would appear, limit our ability to continue this decomposition process. See also John Rogers, ‘Issue Voting’ by Multimember Appellate Courts: A Response to Some Radical Proposals, 49 Vand. L. Rev. 997, 1001 – 1006 (1996) (arguing that the “issues” raised in any case can always be expanded “one level down”).

²⁸Source and executable code for this simulation model are available from the authors.

²⁹There was one additional step that we implemented in our simulation model: If any tree had more than 10,000,000 “leaves,” the simulation for that tree was terminated. This was simply for computational convenience, as keeping track of trees of this size began to exceed the memory and processor limitations of the computer on which the simulation was running.

³⁰We recognize that this simple model does not adequately capture the many nuances of the actual tree-generating process in legal opinions. The way that the legal questions in a real case are expanded and the place(s) where that expansion terminates – the shape and structure of any individual case’s final branching tree – are determined by an extremely complex interaction among a large number of factors: the substantive law of copyright, the facts of the case, the skill of the opposing lawyers, the judge’s expertise, the amount of time and energy that the judge (or her law

clerks) can devote to the matter, the parties' ability to pay for continued downward expansion of the issues in the case, among others. Our model incorporates radical ignorance of these various factors or the way that they interact; we assume simply that the results of a simulated coin toss determine whether or not the expansion process continues or terminates.

³¹ The distributions in Figures 8 – 10 are plotted against double logarithmic scales, for reasons that will become clearer in a moment.

³² The expected value of the number of branches added at any single end-point of this model is simply $p * B$; in the simulated populations shown in Figure 8, this expected value was less than 1.

³³ On double logarithmic scales, an exponential distribution will have a concave shape, as shown by the least-squares regression lines for the exponential distributions plotted in Figure 8.

³⁴ There is some finite probability — equal to p^n — that any individual tree continues to add branches for n recursions, where n is any arbitrarily selected number. But the likelihood that it will do so declines exponentially, and the rapidity of exponential decay insures that this probability becomes vanishingly small in relatively short order. If $p = 0.5$, for example, the probability that any branch will continue for 1000 recursions is $(0.5)^{1000} \cong 10^{-31}$, a very small number indeed.

³⁵ We say “apparently unlimited” because, as noted above, see *supra* note 29, we arbitrarily stopped counting when trees reached a certain size — 10,000,000 leaves — because of memory and processor limitations in the computer used for this simulation. We cannot, therefore, say with assurance that the trees in this portion of the simulated populations *never* stopped growing; nor can we eliminate the possibility that some kind of order re-establishes itself in these distributions, many orders of magnitude distant from the first portion of the population.

³⁶ Equation 3 can be rewritten as

$$\log(N(C)) = \log(\text{constant}) + (-x) * (\log(c)).$$

Power law distributed data will thus plot as a straight line (with slope $-x$) on double logarithmic scales.

³⁷That is, if $N(c_1) = K_1 c_1^{-x}$, then $N(c_2) = K_2 c_2^{-x}$ for all c_2 ; the proportionality constant K changes, but the underlying proportionality relationship does not. The change of scale does not modify the basic statistical behavior of the function. In Schroeder's words:

“Think of a homogeneous power function such as $f(x) = cx^{-a}$ where c and a are constants. . . . Such simple power laws, which abound in nature, are in fact self-similar; if x is rescaled (multiplied by a constant) then $f(x)$ is still proportional to x^{-a} , albeit with a different constant of proportionality. [P]ower laws, with integer or fractional exponents, are one of the most fertile fields and abundant sources of self-similarity.”

Schroeder, *supra* note 6, at 103.

³⁸See, for example, Kauffman and Johnsen's model of the growth of co-evolving biological species within ecosystems. Stuart Kauffman and Sonke Johnsen, "Coevolution to the edge of chaos: Completed fitness landscapes, poised states, and coevolutionary avalanches," 149 *J. Theor. Biol.* 467, 489-91 (1991). Depending on the initial parameter settings of these models, populations of species either “freeze” with little evolutionary change — what Kauffman and Johnsen refer to as the “solid” state — or they fluctuate chaotically — the “gas” state. And “[j]ust as at the interface between the solid and gas phase is a kind of 'liquid' region,” so too is there an interface or boundary where evolutionary change proceeds in accordance with a power law.” *Id.* Similarly, the simple logistic model of population growth, $N(c') = (1-x) * N(c)$ also shows this boundary phenomenon, with narrow regions of self-similar fractal output in between

regions of stable and chaotic growth. See Williams, *supra* note 6, at 161-173, 189-197; Alligood *et al.*, *supra* note 6, at 18-24.

Note that this boundary phenomenon — a power law distribution at the interface between an orderly and a disorderly state — is embodied in the definition of a fractal object as well. As noted above, see *supra* note 10, for true fractals the expression

$$L(\varepsilon) = N(\varepsilon) * \varepsilon^D$$

converges to zero for all $D > D_{\text{crit}}$, diverges to infinity for all $D < D_{\text{crit}}$, and converges to a finite non-zero number at D_{crit} (the fractal dimension of the object). The fractal dimension is thus, in this sense, a unique number at the “boundary” between regions where the length function $L(\varepsilon)$ “explodes” ($D < D_{\text{crit}}$) or disappears ($D > D_{\text{crit}}$).

³⁹The size of earthquakes (the “Gutenberg-Richter law”), see Per Bak, *How Nature Works* (1996) at 12 - 14, the size of cities in the United States (and elsewhere), see Paul Krugman, *Development, Geography, and Economic Theory* (1997) 42 - 46, Schroeder, *supra* note 6, at 103-104, the size of meteorites, see A.Z. Mekjian, Model of a fragmentation process and its power-law behavior, 64 *Phys. Rev. Lett.* 2125 (1990), the size of biological organisms (“Cope’s Law”), see Schroeder, *supra* note 6, at 105, R.M. May, How many species are there on earth?, 214 *Science* 1441 (1988), the size and shape of rainforest patches and the growth of different species populations within those patches, see G. Sugihara and R.M. May, Applications of fractals in ecology, 5 *Trends Ecol. Evol.* 79, 82-3 (1990), R.V. Sole and S.C. Manrubia, "Are rainforests self-organized in a critical state?" 173 *J. Theor Biol* 31 (1995), the size and shape of clouds, see Schroeder, *supra* note 6, at 221, the growth of blood vessels and neurons, see *id.*, at 119, the size of personal incomes (“Pareto’s law”), see *id.*, at 34-5, the number of species becoming extinct in

any period, see Ricard V. Sole, Susanna C. Manrubia, Michael Benton, and Per Bak, Self-similarity of extinction statistics in the fossil record, 388 *Nature* 764 (1997), Kauffman and Johnsen, *supra* note 35, at 493, Bak, *supra*, at 164-65, the number of times particular words occur in spoken or written text (“Zipf’s Law”), see Schroeder, *supra* note 6, at 35, Bak, *supra*, at 24-5, the growth of traffic jams, Bak, *supra*, at 197, even the growth in the number of scientific journals devoted to particular subjects over time (“Bradford’s Law”), see S. Naranan, Bradford’s Law of Bibliography of Science: An Interpretation, 227 *Nature* 631 (1970), Robert Fairthorne, Empirical Hyperbolic Distributions for Bibliometric Description and Prediction, 25 *J. Doc* 319 (1969)— all appear to be distributed in accordance with a power law function.

This rather extraordinary regularity, long written off as some sort of peculiar statistical oddity, has begun to attract increasing attention; it is likely that this statistical regularity reflects some important regularity in the underlying processes at work in these diverse phenomena. See, for example, Paul Krugman, *Development, Geography, and Economic Theory* (1997) at 42 - 46 (putting forth the “ghostly beginnings of an explanation” for the “spooky regularity” of the power law distribution for city size, and noting that the distribution is “a major embarrassment for economic theory; one of the strongest statistical relationships we know, lacking any clear basis in theory”); Bak, *supra*, at 14 (“The importance of the Gutenberg-Richter law [for earthquake magnitudes] cannot be exaggerated. It is precisely the observation of such simple empirical laws in nature that motivates us to search for a theory of complexity) (emphasis in original); M.E.J. Newman and Gunther J. Eble, *Decline in extinction rates and self-similarity in the fossil record*, 1998 Santa Fe Institute Working Papers Series Number 98-04-03, at 2 (ubiquity of power law

distributions in many branches of science and engineering “poses one of the great unsolved scientific mysteries of the last few decades”).

⁴⁰The theory of “self-organized criticality” was first advanced in P. Bak, C. Tang, and K. Wiesenfeld, *Self-organized Criticality: an Explanation for 1/f Noise*, 59 *Phys. Rev. Lett.* 381 (1989), and was further amplified (in a non-technical manner) in Bak, *supra* note 36 (developing an extensive argument in favor of theories of self-organized criticality in complex systems of all kinds). See, generally, Schroeder, *supra* note 6, at 389; C. Adami, *Self-organized Criticality in Living Systems*, 203 *Phys. Lett. A* 29-32 (1995) (self-organized criticality refers to systems “that are driven to a critical state which is robust to perturbations and whose macroscopic behavior is predictable to the extent that it follows power laws with exponents depending on geometry and spatial structure”); R.V. Sole and S.C. Manrubia, *Are Rainforests Self-organized in a Critical State?*, 173 *J. Theor Biol* 31 (1995).

⁴¹William M. Landes and Richard Posner, *supra* note 4, at 252.

⁴²There was no particular method to this selection, other than a suspicion that there might be some temporal diversity in these data, or some differences between citation patterns in federal and state courts, that might be worth exploring. The number of reported opinions in each cohort is shown in Table 1.

⁴³The distributions in Figures 11 – 14 are, again, plotted against double logarithmic scales; power law distributions will display as straight lines on such scales. See *supra* note 36.

⁴⁴In no case, however, was an exponential distribution a *more* effective predictor of the data than the power law distribution, that is, in no case did an exponential function explain a greater fraction

of the variance in the data than did a power law function. In a number of cases, however, the difference between the two functions was not statistically significant.

⁴⁵ Quite apart from the validity of our fractal structure hypothesis as a tool for explaining patterns in these data, certain of those patterns are puzzling and should be the subject of additional work, most notably the marked differences between citation patterns in the federal and state courts that we included in our analyses and that we currently regard as inexplicable. Landes and Posner found similarly puzzling differences between the “depreciation rates” for federal and state common-law cases, see William M. Landes and Richard A. Posner, *Legal Change, Judicial Behavior, and the Diversity Jurisdiction*, 9 *J. Legal Studies* 367, 378-86 (1980). and we suspect that a fuller understanding of these differences can help illuminate more generally the important role of case citations within the legal system.

⁴⁶ Lon Fuller, *Adjudication and the Rule of Law*, 1960 *Proc. Am Society Int'l Law* 1, 8 (1960).

⁴⁷ See especially Stuart A. Kauffman, *Origins of Order* (1993), which remains the best summary of the general theory of complex systems.

⁴⁸ There is a small, but growing, literature on the study of the law as a complex system. See, for example, David G. Post and David R. Johnson, ‘Chaos Prevailing on Every Continent’: Towards a New Theory of Decentralized Decision-Making in Complex Systems, *Chicago-Kent L. Rev.* (in press); Michael S. Fried, *The Evolution of Legal Concepts: The Memetic Perspective*, 39 *Jurimetrics* 291 (1999); J.B. Ruhl, *The Fitness of Law: Using Complexity Theory to Describe the Evolution of Law and Society and Its Practical Meaning for Democracy*, 49 *Vand. L. Rev.* 1407 (1996); J.B. Ruhl, *Thinking of Environmental Law as a Complex Adaptive System: How to Clean Up the Environment by Making a Mess of Environmental Law*, 34 *Houston L. Rev.* 933

(1997); J.B. Ruhl and Harold J. Ruhl, Jr., *The Arrow of the Law in Modern Administrative States: Using Complexity Theory to Reveal the Diminishing Returns and Increasing Risks the Burgeoning of Law Poses to Society*, 30 U.C. Davis L. Rev. 405 (1997); Mark J. Roe, *Chaos and Evolution in Law and Economics*, 109 Harv. L. Rev. 641 (1996); Thomas Geu, *The Tao of Jurisprudence: Chaos, Brain Science, Synchronicity, and the Law*, 61 Tenn. L. Rev. 933 (1994); Andrew W. Hayes, *An Introduction to Chaos and Law*, 60 UMKC L. Rev. 751 (1992).