

Does Age (Really) Matter? A Response to Manning, Carroll, and Carp*

Lee Epstein, *Washington University*

Andrew D. Martin, *Washington University*

For some months now, we have been at work on a project called “Grouping Data.” In it, we address several specific questions (e.g., Under what circumstances should we group data? If the appropriate circumstances exist, how ought we proceed?). But our overarching message is a simple one: the decisions we make over how to group data may crucially affect the inferences we reach.

This certainly holds for “Does Age Matter?” In it, Manning, Carroll, and Carp group a continuous variable, age, into three categories or cohorts: 1 = under the age of 46; 2 = 46–64; and 3 = 65+. When they enter this (now ordinal-level, treated as if it were interval-level) variable into a logistic regression model designed to explain judicial decisions in age discrimination cases (coded 0 if the judge ruled against the plaintiff alleging age discrimination and 1 if in favor of the plaintiff), it produces a positive and statistically significant coefficient; when they enter this same cohort-age variable into a model designed to explain outcomes in litigation alleging discrimination based on race or gender, it yields an insignificant coefficient. Taken together, these results lead the authors to conclude that age “matters” in precisely those cases in which they expect it to matter—age discrimination suits—and that it does not matter in those cases in which they expect it not to matter—race and gender discrimination suits.

As it turns out, though, this conclusion—and, more important, the inference that Manning and his colleagues reach based on their sample of data—rests entirely on their grouping decision. Specifically, if we respecify the models in precisely the same way as the authors but vary the cutoff

*Direct correspondence to Lee Epstein, Edward Mallinckrodt Distinguished University Professor of Political Science and Professor of Law, Washington University, Campus Box 1063, St. Louis, MO 63130 (epstein@artsci.wustl.edu). Andrew D. Martin is Assistant Professor of Political Science at Washington University in St. Louis. We thank Robert Carp, Kenneth Manning, and Bruce Carroll for supplying us with their data. We also are grateful to the members of Washington University’s weekly Workshop on Empirical Research and the Law (WERL) for providing us with insightful comments and suggestions. All data and documentation necessary to replicate the results we present here are located at (<http://www.artsci.wustl.edu/~polisci/epstein/>). We used R (<http://www.r-project.org>) to analyze the data.

points on the age-cohort variable, we find that for age discrimination cases, only 2 of 23 (plausible) cutoff points produce results in line with the authors' expectation.

These findings may bring into question the authors' inferences about the importance (or lack thereof) of age on their own terms. But because their "own terms" rely on some strong (and unnecessary) assumptions,¹ we thought it prudent to properly specify the model, creating two indicator variables to represent the age cohorts rather than entering a single variable with three categories. The outcomes from this set of analyses are even starker: only when we code the age variables exactly as the Manning team has done (i.e., select its chosen cutoff points) does the effect of age run in the hypothesized direction; no other grouping choices produce the expected findings.²

In what follows, we elaborate on these results. We also present analyses of the effect of age on judicial decisions in race and gender cases.

Age and Judicial Outcomes in Age Discrimination Litigation: The Manning et al. Specification

"Does Age Matter?," as we note above, slices the continuous variable of age into three cohorts: judges 45 years old and younger (Cohort 1), those 46–64 years old (Cohort 2), and judges 65 years and older (Cohort 3). Since 65 is the traditional age of retirement, the Cohort 3 cutoff value makes reasonable substantive sense. The value of 45, however, seems arbitrary—and the authors' explanation for it not particularly helpful in justifying or clarifying their decision.

In light of the lack of a theoretically or intuitively pleasing account for the "45" grouping decision, we thought it prudent to determine just how sensitive the Manning et al. results were to it. Accordingly, we fit the three-cohort model for age discrimination cases for all cutoff points between Cohort 1 and Cohort 2 ranging from 40–62;³ we left the third cohort at

¹See the section "The Three-Cohort Model Using a Proper Specification" for more details.

²Our purpose here is to demonstrate how grouping decisions may affect inference; it is not to account for the results of our analyses. Nonetheless, it is probably true that our findings will not come as much of a surprise to judicial specialists. After all, for over a decade now they have attempted to explore whether male and female and black and white judges reach distinct decisions—with their findings mixed at best (for recent reviews, see Epstein and Mather, 2003; Schneider, 2001). If race and gender, typically stable characteristics, do not (depending on the study) significantly affect judicial decisions, then we might hypothesize an even lesser role for aging (which we all, with any luck, undergo). This may be especially so for litigation involving the Age Discrimination in Employment Act, which not only protects "the elderly" (unless we define 40 as "elderly") from employment discrimination based on age, but also covers individuals who are 40 years of age or older.

³Moving the cutoff point to lower than 40 or higher than 62 left too few judges in either Cohort 1 or 2 to make analyses sensible.

TABLE 1

Z-scores (Coefficients Divided by Standard Errors) for All Explanatory Variables in the Authors' Model A (Table 2) Using All Feasible Cutoff Values

Cutoff	Intercept	Age Cohort	Party	ADEA
40	-3.448	1.437	2.055	2.030
41	-3.632	1.608	2.021	2.062
42	-3.914	1.871	1.983	2.098
43	-4.129	2.063	1.979	2.096
44	-4.049	1.877	2.006	2.062
45	-4.496	2.324	1.929	2.135
46	-3.889	1.454	2.072	1.997
47	-4.145	1.703	2.014	2.051
48	-4.017	1.379	2.080	1.972
49	-3.798	1.022	2.153	1.923
50	-3.638	0.711	2.216	1.863
51	-3.459	0.403	2.279	1.817
52	-3.692	0.574	2.238	1.850
53	-3.553	0.306	2.277	1.804
54	-3.947	0.689	2.197	1.873
55	-3.792	0.401	2.271	1.821
56	-3.693	0.179	2.324	1.777
57	-3.802	0.200	2.315	1.782
58	-3.406	-0.469	2.461	1.661
59	-3.547	-0.366	2.437	1.671
60	-3.852	-0.037	2.353	1.731
61	-4.308	0.514	2.230	1.845
62	-4.555	0.758	2.165	1.891

NOTE: The first age cohort is defined as those judges less than or equal to the cutoff; in the second are those older than the cutoff and younger than 65; and the third contains those judges age 65 and older. The boldfaced row is the one reported by the authors.

65.⁴ In all other ways, we followed the authors' approach: we included the two other variables they found to be significant determinants of age discrimination decisions (political party and the Age Discrimination in Employment Act (ADEA)) (see Model A in their Table 2) and we used logistic regression, estimated via maximum likelihood.

Table 1 presents the results, in particular the Z-scores for all feasible cutoff points. If a score falls below -1.96 , then we can conclude that age exerts a statistically significant negative effect (for level of significance $\alpha = 0.05$) on judicial decisions in age discrimination litigation; if it falls above 1.96 , then we can conclude the variable exerts a statistically significant positive effect. If

⁴We want to reiterate that this approach to incorporating age in the model—as a variable with three categories rather than as two indicator variables—requires some very strong assumptions that we (and probably the authors themselves) would prefer not to make. Nonetheless, because we think it important to reproduce their analysis, we follow their approach here; in the next section, we incorporate age in a more appropriate fashion.

TABLE 2

Estimated Coefficients for All Explanatory Variables in the Authors' Model A
(Table 2) Using All Feasible Cutoff Values

Cutoff	Intercept	Age Cohort	Party	ADEA
40	-1.627	0.275	0.388	0.401
41	-1.691	0.303	0.382	0.407
42	-1.782	0.344	0.374	0.414
43	-1.839	0.371	0.373	0.412
44	-1.724	0.324	0.378	0.405
45	-1.867	0.389	0.364	0.420
46	-1.496	0.227	0.390	0.391
47	-1.568	0.260	0.380	0.403
48	-1.423	0.199	0.392	0.385
49	-1.300	0.143	0.406	0.376
50	-1.198	0.096	0.418	0.363
51	-1.108	0.053	0.430	0.354
52	-1.147	0.073	0.423	0.362
53	-1.075	0.038	0.432	0.353
54	-1.164	0.084	0.416	0.367
55	-1.092	0.048	0.429	0.357
56	-1.039	0.021	0.439	0.349
57	-1.041	0.023	0.438	0.350
58	-0.902	-0.053	0.466	0.325
59	-0.925	-0.041	0.462	0.328
60	-0.991	-0.004	0.448	0.340
61	-1.095	0.056	0.423	0.363
62	-1.136	0.082	0.411	0.372

NOTE: The first age cohort is defined as those judges less than or equal to the cutoff; in the second are those older than the cutoff and younger than 65; and the third contains those judges age 65 and older. The boldfaced row is the one reported by the authors.


the score falls between these values, then we cannot reject the null hypothesis of “no effect.”

Notice that we were able to replicate the authors' results: when age is set at their cutoff points, it—along with party and ADEA—significantly affects judicial decisions in the hypothesized direction (i.e., older judges are more likely to rule in favor of plaintiffs in age cases). But—and this is a big but—while the party and ADEA variables show stable effects across all specifications, the age-cohort variable does not. In fact, it is *only statistically significant in 2 of the 23 possibilities*. At the conventional level of significance ($\alpha = 0.05$), we would expect to find a significant relationship in just over one possibility due solely to chance alone.

Figure 1 and Table 2 reinforce these results. In the former we plot the absolute value of the Z-scores, with the horizontal line set at 1.96. As we can see, the findings are simply not robust. In the latter, provided for completeness, we present the estimated coefficients for all three explanatory variables. Again, while the magnitudes of the party and ADEA variables

FIGURE 1

Absolute Values of Z-Scores (Coefficients Divided by Standard Errors) for the Age-Cohort Variable in the Authors' Model A (Table 2) Using All Feasible Cutoff Values



NOTE: The first age cohort is defined as those judges less than or equal to the cutoff; in the second are those older than the cutoff and younger than 65; and the third contains those judges age 65 and older. Forty-five is the cutoff used by the authors.

remain steady across the specifications, the impact of age cohort varies considerably.

The Three-Cohort Model Using a Proper Specification

The approach we have taken thus far to the age-cohort variable corresponds to the authors', and can be represented as follows:

$$\mu_i = \beta_0 + \beta_1 \text{Age Cohort}_i + \beta_2 \text{Party}_i + \beta_3 \text{ADEA}_i \quad (1)$$

Characterizing the variable in this way implies that the difference in the linear predictor (as well as the predicted probability) between Cohorts 1 and 2 is *precisely the same* as the difference between Cohorts 2 and 3. To be sure, this may turn out to be the case *ex post* but it is an unnecessary (and unnecessarily strong) assumption to make at the onset;⁵ it may even be an

⁵From our perspective there seem to be two model-based justifications for grouping data in this fashion. One is a strong theory that prescribes precise cutoff values and cohort effects; the other is when estimating an additional parameter is too costly due to an extremely small sample size.

unwise one for this particular project given Manning et al.'s claims in their footnote 5.⁶ A far more preferable initial strategy is to rewrite the model with two indicator variables:

$$\mu_i = \beta_0 + \beta_1 \text{Age Cohort } 2_i + \beta_2 \text{Age Cohort } 3_i + \beta_3 \text{Party}_i + \beta_4 \text{ADEA}_i \quad (2)$$

This specification leaves the first age cohort (those judges less than or equal to the cutoff value) as the baseline category. Of course, the choice of that baseline is completely arbitrary but *in no way does it affect the findings*. We also should note that if the estimated coefficient for Cohort 3 is twice the size of the coefficient for Cohort 2, then the authors' strategy (Equation (1)) is reasonable. Otherwise, the respecification above (Equation (2)) is preferred.

As it turns out, the estimated coefficients support the wisdom of the two-indicator approach (see Table 4). Perhaps even more interesting, though, are the Z-scores for all variables in the respecified model, which Table 3 depicts. Note that for the cutoff value invoked in the Manning et al. paper, significant effects emerge for the two cohorts: both are more likely to support the plaintiff in age discrimination cases. But that is the only cutoff where both cohort effects are significant. If the cutoff is between 56 and 62, the model suggests that *those in Cohort 2 are less likely than those in Cohort 1 to support plaintiffs in age discrimination cases. Only when the cutoff is set at the authors' choice of 45 does the effect run in the hypothesized direction*. The results are similar for Cohort 3: only when we establish the cutoff at 43 or 45 do the judges exhibit a statistically significant difference from the baseline (Cohort 1).

For completeness and purposes of comparison (with Table 2 and Figure 1) we present the estimated coefficients of all explanatory variables (in Table 4) and the absolute values of the Z-scores for the age variables (in Figure 2). The results parallel those we obtained using the authors' specification: while the coefficients for the party and ADEA variables remain relatively constant across specifications, the age-cohort variables do not.

⁶Specifically, the authors offer the following explanation for their choice of grouping age into cohorts rather than treating it as a continuous variable: "[W]e don't anticipate that the effect of age will be perfectly linear along the age spectrum. There is no theoretical justification to believe that as a judge ages by one year we would anticipate that they would exhibit a correspondingly incremental increase in support for older litigants. We anticipate that differences will emerge among the very oldest and youngest jurists." In light of this claim, it seems particularly ill advised to expect the difference between Cohorts 1 and 2 and 2 and 3 to be precisely the same.

As an aside, Manning and his colleagues (correctly) note that age, when left as a continuous variable, is not a significant predictor of judicial decisions in age discrimination cases (see their footnote 18). The full results are as follows:

	Estimate	SE	Z Value	Pr(> z)
(Intercept)	-0.8831	0.6453	-1.37	0.1711
Age	-0.0020	0.0106	-0.18	0.8533
Party	0.4539	0.1892	2.40	0.0165
ADEA	0.3346	0.1963	1.70	0.0883

TABLE 3

Z-Scores (Coefficients Divided by Standard Errors) for All Explanatory Variables in the Authors' Model A (Table 2) Using All Feasible Cutoff Values and the Specification in Equation (2)

Cutoff	Intercept	Age Cohort 2	Age Cohort 3	Party	ADEA
40	-2.172	0.085	0.683	1.995	2.060
41	-2.675	0.551	1.125	2.003	2.059
42	-3.323	1.178	1.703	2.029	2.031
43	-3.742	1.592	2.073	2.091	1.931
44	-3.808	1.152	1.766	2.035	1.985
45	-4.585	2.011	2.475	2.071	1.940
46	-3.920	0.340	1.258	1.992	2.049
47	-4.356	0.806	1.616	1.982	2.052
48	-4.344	0.239	1.281	1.988	2.049
49	-4.138	-0.371	0.936	2.002	2.066
50	-4.047	-0.906	0.666	2.018	2.101
51	-3.915	-1.432	0.418	2.037	2.131
52	-4.264	-1.131	0.678	2.037	2.063
53	-4.198	-1.601	0.486	2.111	2.049
54	-4.678	-0.972	0.874	2.049	2.053
55	-4.541	-1.527	0.692	2.056	2.032
56	-4.453	-2.015	0.579	2.051	1.978
57	-4.640	-2.088	0.643	2.066	2.002
58	-4.410	-3.541	0.256	2.136	2.074
59	-4.554	-3.585	0.363	2.158	2.010
60	-4.857	-3.215	0.571	2.236	1.952
61	-5.247	-2.310	0.948	2.099	1.990
62	-5.567	-2.021	1.109	2.110	2.096

NOTE: The first age cohort is defined as those judges less than or equal to the cutoff; in the second are those older than the cutoff and younger than 65; and the third contains those judges age 65 and older. The boldfaced row is the one reported by the authors.

Taken collectively, our analyses so far suggest an effect of age only if we are willing to accept (for whatever reason) the Manning et al. cutoff of 45. Let us assume, for the sake of argument, that we could, say, develop a plausible story to support the authors' grouping decision. What we still must determine is whether the estimated coefficients for Cohort 2 and Cohort 3 differ significantly in the hypothesized direction (i.e., the effect of Cohort 3 should be significantly greater than that of Cohort 2). We can do so by testing the null hypothesis that $\beta_2 = \beta_3$ using a likelihood ratio test. The test statistic is 1.165, which is distributed as a χ^2 random variable with one degree of freedom. The associated p -value is 0.720—meaning that there is no statistically significant difference between the estimated coefficient for Cohort 2 and Cohort 3; whatever difference exists is likely due to chance alone. This, in turn, suggests that a single threshold exists at age 45; there are no differences between those judges above and below age 45, given this cutoff.

TABLE 4

Estimated Coefficients for All Explanatory Variables in the Authors' Model A (Table 2) Using All Feasible Cutoff Values and the Specification in Equation (2)

Cutoff	Intercept	Age Cohort 2	Age Cohort 3	Party	ADEA
40	-1.143	0.044	0.365	0.379	0.408
41	-1.362	0.274	0.583	0.381	0.408
42	-1.633	0.566	0.854	0.386	0.403
43	-1.796	0.753	1.018	0.400	0.384
44	-1.497	0.443	0.721	0.388	0.394
45	-1.764	0.752	0.987	0.396	0.386
46	-1.182	0.098	0.406	0.378	0.406
47	-1.285	0.226	0.509	0.377	0.407
48	-1.146	0.061	0.370	0.378	0.406
49	-1.038	-0.089	0.258	0.380	0.409
50	-0.963	-0.210	0.177	0.384	0.417
51	-0.899	-0.323	0.108	0.388	0.423
52	-0.954	-0.247	0.171	0.388	0.409
53	-0.911	-0.345	0.120	0.403	0.407
54	-0.995	-0.208	0.211	0.390	0.407
55	-0.947	-0.327	0.164	0.392	0.403
56	-0.912	-0.441	0.134	0.391	0.393
57	-0.928	-0.471	0.147	0.394	0.398
58	-0.858	-0.873	0.057	0.412	0.415
59	-0.877	-0.965	0.081	0.416	0.402
60	-0.925	-0.956	0.126	0.431	0.390
61	-0.992	-0.750	0.207	0.401	0.396
62	-1.037	-0.762	0.240	0.403	0.416

NOTE: The first age cohort is defined as those judges less than or equal to the cutoff; in the second are those older than the cutoff and younger than 65; and the third contains those judges age 65 and older. The boldfaced row is the one reported by the authors.


Race, Gender, and Judicial Outcomes in Age Discrimination Litigation

To bolster their claims about the importance of age, Manning and his colleagues make an interesting move: they assert that age should matter only in the resolution of age discrimination litigation, not in the outcomes of suits alleging discrimination based on race and gender. Since they demonstrate the veracity of this claim using the same age-cohort variable and the same specification as they invoked for age discrimination litigation, we thought it appropriate to analyze the impact of their decisions—the chosen cutoff points, as well as the use of one variable to represent age (see Equation (1))—on their results.

For the sake of space, Table 5 presents the Z-scores for the age-cohort variables only—with columns labeled “Original” using the authors’ specification (Equation (1)) and columns labeled “Cohort 2” and “Cohort 3” invoking the more appropriate specification (Equation (2)). Beginning with the outcomes in race discrimination cases, if we use the authors’

FIGURE 2

Absolute Value of Z-Scores (Coefficients Divided by Standard Errors) for the Age-Cohort Variable in the Authors' Model A (Table 2) Using All Feasible Cutoff Values and the Specification in Equation (2)



NOTE: The first age cohort is defined as those judges less than or equal to the cutoff; in the second are those older than the cutoff and younger than 65; and the third contains those judges age 65 and older.

specification, the findings are as they report: the age cohort never exerts a statistically significant effect on the outcomes of these cases. The more appropriate two-variable approach, however, yields a different set of results: for a large number of plausible cutoff values, Cohort 2 judges are significantly less likely to decide in favor of those alleging discrimination

TABLE 5

Z-Scores for Age-Cohort Variables (Coded as a Single Variable (Original) and as Two Dummy Variables (Cohort 2 and Cohort 3) for the Race and Gender Discrimination Models Presented in the Authors' Table 3

Cutoff	Race			Gender		
	Original	Chohort 2	Chohort 3	Original	Chohort 2	Chohort 3
40	-0.102	-0.653	-0.589	1.787	-0.776	0.029
41	-0.635	-2.005	-1.809	1.678	-0.924	-0.002
42	-0.690	-1.891	-1.633	1.663	-0.835	0.175
43	-0.587	-1.517	-1.254	1.459	-1.244	-0.128
44	-0.563	-1.399	-1.120	1.207	-1.666	-0.388
45	-0.468	-1.108	-0.825	0.950	-1.971	-0.449
46	-0.398	-0.938	-0.669	1.007	-1.681	-0.019
47	-0.783	-1.625	-1.146	0.932	-1.699	0.135
48	-0.970	-1.913	-1.283	0.568	-2.250	-0.096
49	-1.416	-2.654	-1.716	0.343	-2.582	-0.183
50	-1.098	-2.072	-1.279	0.250	-2.701	-0.140
51	-1.074	-2.012	-1.171	0.617	-2.028	0.421
52	-1.325	-2.442	-1.371	0.776	-1.741	0.724
53	-0.877	-1.665	-0.860	1.167	-1.061	1.214
54	-0.620	-1.224	-0.573	1.006	-1.367	1.143
55	-0.935	-1.793	-0.817	1.200	-1.059	1.384
56	-1.163	-2.241	-0.952	1.386	-0.762	1.584
57	-1.513	-2.950	-1.165	2.026	0.376	2.120
58	-1.199	-2.444	-0.879	2.197	0.675	2.253
59	-1.308	-2.803	-0.902	2.162	0.577	2.226
60	-1.216	-2.742	-0.819	2.008	0.201	2.119
61	-1.014	-2.582	-0.630	1.928	-0.093	2.066
62	-0.878	-2.631	-0.518	2.001	-0.053	2.104

NOTE: The boldfaced row is the one reported by the authors.

than are those in Cohort 1 (though no significant differences ever emerge between Cohort 3 and Cohort 1 judges).

Turning to the gender discrimination cases, whether we consider the original or more appropriate specification, we observe a number of cutoff values that produce significant results. So, for example, if we group the data between 56 and 63, there is evidence that those judges in Cohort 3 are more likely than those in Cohort 1 to support the result favored by the plaintiff. Even more intriguing is that if we use the authors' cutoff of 45, significant differences too emerge—with Cohort 2 judges less likely to support claims of gender discrimination than their counterparts in Cohort 1.

Discussion

Does age matter? This is the question that Manning and his colleagues pose, and it is one they answer in the affirmative for age discrimination suits and in the negative for others.

Whether they have reached a good inference is the question we pose, and it is one we leave to readers to answer based on the analyses we have presented. For our primary purpose here was not to undermine the authors' study but rather to investigate whether and in what ways their "data grouping" (and, to a lesser extent, their model specification) choices affected their conclusions.

This much we think we have accomplished—with the overall findings taking us back to our starting point: the decisions we make over how to group data may crucially affect the inferences we reach. To see this, we need only consider the answer Manning et al. might have derived had they continued to invoke the single-variable specification of age but set their first cutpoint just one year higher or lower than 45: that age is not significantly related to outcomes in age discrimination cases. On the other hand, had they made different grouping decisions in their analyses of race and gender cases, they might have reported far different results—that, in fact, age "matters" in the resolution of these cases.

In light of the critical role grouping choices play in making inferences, what steps should researchers take to ensure that those they make are of the highest quality? Since we address this in our larger project—the product of which we hope to disseminate shortly—suffice it to note here that we urge scholars to consider, as a threshold matter, whether the circumstances are such that they should be grouping data. If not, they ought avoid it altogether and make use of a linear or, if the project so dictates, a nonlinear specification.⁷ That is because grouping ungrouped data is a procedure that

⁷Given the authors' notions about the relationship between judicial age and outcomes in age discrimination suits (see our footnote 6), we thought it reasonable to fit a model with the natural logarithm of age as a covariate, as well as a quadratic specification with age and age-squared. As was the case when we entered age as a continuous variable (again, see footnote 6), in neither of these specifications was age a statistically significant predictor of outcomes. The results are as follows.

	Estimate	SE	Z Value	Pr(> z)
(Intercept)	-0.7577	2.3957	-0.32	0.7518
ln(Age)	-0.0593	0.5885	-0.10	0.9198
Party	0.4500	0.1885	2.39	0.0170
ADEA	0.3384	0.1954	1.73	0.0833

Logistic regression with the natural logarithm of age as an independent variable.

	Estimate	SE	Z Value	Pr(> z)
(Intercept)	-3.2533	3.1748	-1.02	0.3055
Age	0.0839	0.1190	0.74	0.4580
Age ²	-0.0008	0.0010	-0.76	0.4458
Party	-0.4794	0.1924	2.49	0.0127
ADEA	0.2963	0.2022	1.47	0.1428

Logistic regression with age and age-squared as independent variables.

almost always discards information and thus, as a general matter, a procedure to shun (Epstein and King, 2002). On the other hand, if sufficient reasons exist to group, then we advise researchers to assess systematically the extent to which their conclusions reflect their decisions. Only by conducting the sort of sensitivity analyses that we have done here will investigators (and their readers) have a full appreciation of just how their grouping choices—crucial choices as it turns out—affect their findings and resulting inferences.

REFERENCES

- Epstein, Lee, and Gary King. 2002. "The Rules of Inference." *University of Chicago Law Review* 69:1–133.
- Epstein, Lee, and Lynn Mather. 2003. "Beverly Blair Cook." In N. Maveety, ed., *The Pioneers of Judicial Behavior*. Ann Arbor, Mich.: University of Michigan Press.
- Schneider, Daniel M. 2001. "Empirical Research on Judicial Reasoning: Statutory Interpretation in Federal Tax Cases." *New Mexico Law Review* 31:325.