

Statistics as If Politics Mattered: A Reply to Fowler and Hall

Christopher H. Achen, Princeton University

Larry M. Bartels, Vanderbilt University

Anthony Fowler and Andrew Hall question both the statistical validity and the broader significance of our analysis of the electoral impact of shark attacks along the Jersey Shore in 1916. Setting aside the politics and history of the Progressive period, they focus on methodological considerations, carrying out an extensive set of tests intended to cast all-purpose doubt on our interpretation of the electoral evidence. However, we show that this style of analysis leads them not just into historical misjudgments but into statistical lapses as well. Correcting those missteps, and setting aside the substantial share of their evidence that has no bearing on our argument, leaves our conclusions handsomely supported. In the end, we argue, statistical calculations ignoring the relevant politics and history contribute little to scientific understanding.

We thank Professors Fowler and Hall for their discussion of our work on the electoral effects of the 1916 New Jersey shark attacks (Achen and Bartels 2016, 118–28). They present a brisk critique, setting out possible issues with our analysis. We are grateful for the opportunity to address their concerns, clarifying and extending the argument we presented in our book.

Fowler and Hall offer four lines of criticism. We begin by summarizing them, along with our replies, before setting out our evidence in greater detail.

First, Fowler and Hall's own analysis of 97 fatal shark attacks between 1872 and 2012 provides "little support for the hypothesis that shark attacks decrease support for incumbent presidents or their parties" (2018, in this issue, XXX). However, that analysis has no real bearing on our work. Attributing to us the notion that, in general, "shark attacks influence presidential elections," much less that "irrelevant events generally influence presidential elections" (XXX, XXX), reflects a profound misreading of our argument. As we spelled out, the

1916 attacks were politically relevant because substantial economic losses ensued and the president was explicitly blamed. Neither of those things is true of the typical shark attack; thus, we would not expect it to matter at the polls.

Second, Fowler and Hall's reanalysis of county-level voting patterns in New Jersey in 1916 produces "substantively smaller and statistically weaker" estimates of the impact of shark attacks on electoral support for Woodrow Wilson "under alternative specifications" (2018, XXX). However, they get somewhat smaller shark effects with larger confidence intervals only by omitting or distorting the key factors influencing the 1916 election, resulting in regressions that fit much less well than our own. For example, they show that the apparent effect of the shark attacks is one-third smaller when several coastal counties with little or no shore tourism are treated as if they were affected. Much the same thing happens when the political distinctiveness of New Jersey's machine counties in this era is ignored or when Mayhew's (1986) designation of machine counties in the 1960s is substituted for our designa-

Christopher H. Achen (achen@princeton.edu) is Professor of Politics and the Roger Williams Straus Professor of Social Sciences at Princeton University, Princeton, NJ 08540. Larry M. Bartels (bartels@vanderbilt.edu) is May Werthan Shayne Chair of Public Policy and Social Science at Vanderbilt University, Nashville, TN 37203-5721.

Data and supporting materials necessary to reproduce the numerical results in the article are available in the *JOP* Dataverse (<http://thedata.harvard.edu/dvn/dv/jop>). This article was invited by the editors. It went through editorial review with suggestions from one outside reviewer.

The Journal of Politics, volume 80, number 4. Published online July 31, 2018. <http://dx.doi.org/10.1086/699245>
© 2018 by the Southern Political Science Association. All rights reserved. 0022-3816/2018/8004-00XX\$10.00

000

tion based on contemporary historians' accounts. Even so, their point estimates mostly differ only modestly from ours. The estimates become statistically insignificant only because their ahistorical bad statistical fits inflate standard errors and thus make the t -statistics smaller. We show that a variety of regression models that get the politics right all fit better than Fowler and Hall's. Those models all show a substantively and statistically significant shark effect.

Third, Fowler and Hall's reanalysis of township-level voting patterns within Ocean County in 1916 produces "no evidence that the shark attacks hurt Wilson in New Jersey's beach towns" (2018, XXX). Their argument depends critically on modifying our original data, using an implicit assumption that electoral procedures on the Jersey Shore were the same in the early twentieth century as they are now. However, we show that those procedures have changed and, thus, that Fowler and Hall's revised data set is inaccurate. Once we correct this error, the shark effect returns.

Fourth, Fowler and Hall employ a series of "placebo tests" comparing election outcomes in coastal and noncoastal counties to suggest that "Achen and Bartels's result for New Jersey in 1916 was somewhat likely to arise even if shark attacks have no effect on presidential elections" (2018, XXX). They look at 20 different states with ocean shorelines and 36 different election years. They find that 27% of these comparisons produce "statistically significant" differences in the vote swing from one election to the next in counties bordering the ocean; hence, they argue that other factors besides shark attacks could have produced the marked electoral shift in Jersey Shore counties in 1916. But Fowler and Hall provide no indication of what those other factors might be. In any case, their "placebo" would be better described as a well-known drug. No student of American politics should be surprised to discover that Texas counties along the Gulf of Mexico are different in many ways from those inland or that the politics of New York City and Long Island are distinct from those of upstate New York. Showing that such familiar differences have sometimes produced distinctive electoral shifts is easy enough, but it shows only that many things influence elections. It does nothing to demonstrate that the shark attacks made no difference. The fact that one drug works does not prove that another drug has failed.

Thus, while Fowler and Hall's arguments may seem plausible at first glance, especially for readers unfamiliar with New Jersey politics in the Progressive era, they turn out upon careful reading to be wholly unpersuasive.

We now take up in detail Fowler and Hall's four points. In each case, we briefly summarize our original argument, then address their criticisms, showing why we find their claims unconvincing.

SHARK ATTACKS, 1872–2012

The observation that income losses tend to dent political support for incumbent politicians dates to the nineteenth century. It entered political science with Gosnell and Pearson (1939) and Kramer (1971), whose work was followed by hundreds of other studies demonstrating that pocketbook voting generally has strong and reliable electoral effects, especially in the United States.

As we noted (Achen and Bartels 2016, 92), these effects have been interpreted by prominent scholars "as evidence of the fundamental rationality of American voters and elections." But that interpretation requires that voters respond to economic outcomes that reflect incumbents' good or bad performance while ignoring those that are attributable to circumstances beyond the incumbents' control. Voters who mistake luck for skill "significantly degrade the efficacy of elections as mechanisms for selecting and sanctioning political leaders" (115, 102–8, 329–33).

In most settings, economic and social conditions reflect a complex combination of incumbent leaders' skill and luck. In an attempt to sidestep that complexity, we focused part of our analysis on calamitous events that "rational voters ought to recognize as clearly outside any incumbent's control." We argued that "electoral responses to natural disasters are just particularly illuminating instances of the broader phenomenon of retrospective voting"—illuminating because they might show whether incumbents "pay at the polls for bad times, whether or not objective observers can find a rational basis for blame" (Achen and Bartels 2016, 114, 118).

We spelled out in some detail the unusual features of the 1916 New Jersey attacks that made them relevant for testing our argument. First, the attacks caused very substantial local business losses, with Shore resorts suffering 75% vacancy rates in the midst of their high season. Second, since most scientists before 1916 believed that unprovoked sharks did not attack human beings, the attacks had considerable shock value and generated widespread national (and even international) publicity. Third, President Wilson and prominent members of his administration were nearby, and Wilson was explicitly blamed for not solving the problem—a crucial component of our argument, since "voter's attributions of blame are often crucial in their decisions to punish incumbents" (Achen and Bartels 2016, 138, 140).¹

1. Thus, we did not argue that any natural disaster will reduce the vote for any incumbent, as some recent interpretations of our work have implied (Bodet, Thomas, and Tessier 2016). To the contrary, we emphasized that people's pain is unlikely to be electorally relevant unless politicians or groups formulate it as a political problem, a familiar result from political

Contrary to Fowler and Hall, we certainly never claimed that, in general, “shark attacks influence presidential elections” (2018, XXX). To our knowledge, no one has ever made such a claim. Fowler and Hall gesture at the distinctive features of the 1916 case, noting that “the president had vacationed in the area” and that “the attacks affected the local economy,” and they then search for “similar behavior in other instances” (XXX). But instead of looking for other historical events with the theoretically relevant circumstances, they ignore those considerations and instead provide a “comprehensive” analysis of US shark attacks. How many of the 96 other cases they examine involved substantial economic costs in the affected communities, significant publicity, and explicit attributions of blame to the president or the federal government? We suspect that the vast majority of shark attacks fail those conditions and thus are irrelevant to our argument. If that is right, then most shark attacks should have made little or no difference at the polls.

Indeed, Fowler and Hall’s political-context-free approach fails to turn up much evidence for the impact of generic shark attacks. They conclude that our analysis of the 1916 shark attacks “is a one-off case relying on an incredibly unusual constellation of circumstances and therefore has little to say about democratic accountability and voter competence in general” (Fowler and Hall 2018, XXX). Now of course, highly politicized shark attacks that impose heavy economic costs on a region are rare. If shark attacks were the only possible occasion for voter irrationality, then Fowler and Hall would be right to question their broader significance. However, that is far from the case, and thus the frequency of politically relevant shark attacks is beside the point.

Instead, the theoretically meaningful question is how often voters punish politicians for events that politicians do not control. The 1916 shark attacks provide a compelling example of blind retrospection, but that example is certainly not exhaustive of the broader phenomenon.² Indeed, we cited other cases in which economic fluctuations plausibly unrelated to the performance of incumbent leaders seem to

history studies (Apter 1964; Cantril 1941, chap. 3; Edelman 1964, chap. 2; Gaventa 1980; Smith 1992; Tarrow 1998, chaps. 1, 7; and many others). We showed in two historical cases that when that condition failed and incumbents were not explicitly blamed, electoral retribution failed as well. Not mentioning the prior literature, Fowler and Hall (2018, XXX) describe these two positive confirmations of standard political science findings as “null results.”

2. That is, the proportion of interest is the fraction of elections influenced by irrelevant events, which is substantial and important. Fowler and Hall instead concern themselves with the fraction of irrelevant events that influence elections, which is small and immaterial to our argument.

have had significant electoral effects (Achen and Bartels 2016, 114) and added some analysis of voting patterns in the Great Depression era suggesting that incumbents seem to have been rewarded and punished for ups and downs in the economy unrelated to their own policy preferences (200–211).

We also supplemented our analysis of the shark attacks with a large-scale analysis of the electoral impact of droughts and wet spells over the entire twentieth century (Achen and Bartels 2016, 128–35). The logic of that analysis rests on the judgment that significant droughts and wet spells are likely to cause real distress for large numbers of voters. That seems especially likely in rural areas where farming and ranching are major economic activities. Nonetheless, we recognized that “some droughts will have substantial economic and political impacts and others less so” (132) and reported substantial variability over the course of the twentieth century in the apparent electoral impact of climatic conditions. Notwithstanding that variability, as we expected, voters seem to “punish most of the time” (137). Fowler and Hall do not address that finding or its bearing on the argument they seek to undermine.

Much remains to be learned about the discernment of retrospective voters and the impact of their behavior on political leaders. But scientific progress in this area will require careful investigation of circumstances in which voters would plausibly be expected to feel significant pain or pleasure and then attribute it, reasonably or unreasonably, to incumbent politicians. Fowler and Hall’s “comprehensive” analysis of shark attacks simply fails to address the relevant question. Thus, we turn to the more relevant aspects of their critique, which focus specifically on the cogency of our claims regarding the 1916 shark attacks.

COUNTY-LEVEL ANALYSIS

In our first statistical analysis of the shark effect, we compared Woodrow Wilson’s performance in the 1916 presidential election in the four New Jersey Shore counties hard hit by the summer shark attacks with his performance in the rest of the state. We found that Wilson’s 1916 vote fell by about 3 percentage points in the Shore counties.³ The *t*-ratio for the regression coefficient exceeded 3.0, so that the effect easily achieved statistical significance. A variety of other

3. We controlled throughout for Wilson’s 1912 vote share. In New Jersey and most of the Northeast (although not in much of the rest of the country), Wilson got almost exactly the same three-party vote share in 1912 as his two-party share in 1916 (Achen and Bartels 2016, 121). William Howard Taft and Theodore Roosevelt simply split the Republican vote in 1912, and both men’s votes generally returned to Charles Evans Hughes in 1916. Thus, Wilson’s 1912 vote share works very well as a baseline.

statistical specifications led to the same result—a drop of 2–4 percentage points in Wilson’s vote share in the Jersey Shore counties (Achen and Bartels 2016, 122).

Fowler and Hall criticize this analysis on the grounds that it involves “several important specification choices” that are “consequential for the substantive size and statistical strength of the estimated effect” (2018, XXX). They proceed to examine five alternative specifications and, more briefly, hundreds of additional alternatives generated by mechanical mixing and matching of these and other possibilities. In every instance they find some evidence of a shark effect, but in most cases the estimates are smaller and less precise than in our original analysis. If these alternative specifications were “defensible” and “reasonable,” as Fowler and Hall assert (XXX, XXX), then it would be perfectly appropriate to weigh their results along with ours (Bartels 1997). As it turns out, the effect would be too muddy but not substantially alter our conclusions. However, analyses blind to historical and political context will seldom be “defensible” or “reasonable,” and we will show in some detail that that is true of their plethora of “forking paths.”

In our study, we dropped Essex County from our analysis because of the dramatic political circumstances that distinguished it from the rest of the state. Fowler and Hall speculate that perhaps Essex County belongs in the regression after all. They write that how we “decided that Essex was an anomaly while the other 20 counties were worth inclusion is unclear.” They seem to have missed the paragraph on the page facing the statistical results that they are criticizing, where we discussed exactly this point (Achen and Bartels 2016, 122).

After running for governor in 1910 with support from the state Democratic political machines, Wilson had double-crossed the bosses in several ways, including his successful enactment of the Geran Law, which helped clean up corrupt election administration in New Jersey and made Wilson’s national reputation (Ludington 1911). In the book, we cited two prominent Wilson biographers who noted that Jim Smith, Democratic Party boss of Essex, was bitterly opposed to Wilson from 1911 onward and worked to deny him the presidency in 1912; no other county Democratic leaders did the same (Blum 1951, 39–40; Link 1947, 236–37, 288–89, 424; see also Hofstadter 1955, 183 n. 2). Observers in 1912 were well aware that Essex was the only county in the state whose Democratic Party organization was not under Wilson’s control. Speaking of the 1911 elections, Hosford (1912, 19) wrote: “Of course every one [*sic*] in the United States, who reads the newspapers, knows by this time why the old Smith-Nugent machine, of New Jersey, would never again consciously render any more favors to Woodrow Wilson. . . . Humiliated by

their disappointment in not being able to control Governor Wilson and the last Legislature, only one course of action was open to them. Revenge, sweet, delicious revenge! Revenge at any price! . . . It [Essex] was the only county in the State where the Governor did not speak during the campaign.” Subsequent Wilson biographers have followed Blum, Link, and Hosford in noting that, unlike the Democratic parties in the rest of the counties, Essex County opposed Wilson throughout the 1912 presidential campaign (Berg 2013, 229; Cooper 2009, 153). By 1916, however, Wilson’s control of the entire state had slipped away (Blum 1951, 76; Greenberg 2012, 225), making Essex no longer an outlier. The consequence is that 1912 is a good control for every county except Essex.⁴

It is no surprise, then, that if one forces Essex into the regression anyway, as Fowler and Hall do, the unexplained variance doubles, and all the standard errors increase concomitantly—classic symptoms of having included an inappropriate outlier. Indeed, a standard test of whether an observation already under suspicion on substantive grounds does not belong in a regression (the “externally studentized residual test”; see Weisberg 1985, 116) decisively rejects Fowler and Hall’s suggestion that Essex County belongs in the regression ($p < .002$).

Now, an interest in New Jersey political history is a specialized taste, like a fondness for anchovies or stinky tofu. However, even if one knew nothing about the relevant history, the statistical alarm bells set off by including Essex County in a regression of this sort should be sufficient to deter any data analyst from proceeding down this forking path. The result of including Essex is an enormous residual. Even on a very conservative test that assumes substantive ignorance about the observations, treats them all as equally suspicious, and thus leans toward keeping errant observations, the residual is sufficiently large to reject at the .05 level the hypothesis that Essex County belongs in the regression (Weisberg 1985, 116–17 and table E, 301–5). Thus, neither sensible historical judgment nor conservative statistical practice supports proceeding as Fowler and Hall do. Their alternative specification is decisively rejected on both qualitative and quantitative grounds.

If Essex is included in the regression anyway, the estimated shark effect is actually somewhat larger than before, and it remains significant at the .05 level, as we were well aware and as Fowler and Hall’s table 3 reports. Thus, in this particular regression the historical and statistical mistake of

4. For the 1912 presidential election generally, see Gould (2008) and Link (1947); for 1916, see Gould (2016), Link (1965), Lovell (1980), and Nathan (1965).

including Essex makes little difference. But we have gone over this point in some detail because it illustrates clearly the pitfalls of Fowler and Hall's approach. In their critique, the historical context, both what we wrote and what others have written, goes unmentioned or even explicitly disregarded, as if qualitative evidence were irrelevant to quantitative researchers. That is not a recipe for good social science, and especially not for good historical research. Respect for qualitative evidence prevents many a quantitative blunder. As we will see, Fowler and Hall's arguments demonstrate over and over again how important it is to understand historical and cultural contexts and how badly quantitative researchers can go wrong when that step is skipped.

Fowler and Hall next try looking just at the counties in which the shark attacks took place rather than all beach counties. But this approach ignores the historical evidence of economic losses and intense concern about the sharks along the entire Jersey Shore.⁵ Even more oddly, they also try counting as "beach" all New Jersey coastal counties (including Hudson County, whose gritty beachside docks and warehouses provide the backdrop for the Marlon Brando film *On the Waterfront*, conveying a clear sense of just how much family beach tourism might be expected there). These alternative specifications fly in the face of both common sense and standard accounts of the distinctive dependence on tourism of exactly the four Jersey Shore counties we identified (e.g., Edelstein 1999; Roberts and Youmans 1993; Wilson 1953). But in spite of all that, the indomitable shark effect in these implausible regressions remains between 2 and 3 percentage points, one statistically significant and the other nearly so.

Finally, Fowler and Hall question our "machine county" control variable. They acknowledge that "populous counties with large immigrant populations were different from the rest of the state and may have differentially shifted their support for Wilson between 1912 and 1916" (Fowler and Hall 2018, XXX). However, they then proceed to ignore that fact in their own analyses, first excluding the control for ma-

chine counties altogether and then substituting Mayhew's (1986) list of machine counties a half century later on the grounds that "the county machines he identifies are likely to have been long-lived" (Fowler and Hall 2018, XXX). Even a cursory reading of the relevant scholarly literature could have alerted them to significant changes in the number and character of urban machines over that half century (e.g., Banfield and Wilson 1963, 116), while a glance at the demographic characteristics of Mayhew's machine counties would have shown that some of them were not "populous counties with large immigrant populations" in 1916. Again Fowler and Hall's "forking paths" do not get the politics nearly right. Thus, it should not be surprising that both these specifications produce dramatically lower R^2 statistics (although the estimated shark effects still exceed 2 percentage points with t -ratios of 1.2 and 1.5). The obvious question—why do these regressions fit so poorly?—is not one that Fowler and Hall pause to ask.

Now, the fact that Fowler and Hall's treatment of "machine counties" is unpersuasive does not imply that our own approach is the only plausible one. Historical judgments are never beyond dispute, and the construction of this variable required careful assessment of the politics of Wilson's time. Thus, we lay out here the considerations that led us to define it as we did, as well as the implications for our argument of some plausible alternative approaches.

The central political cleavage in early twentieth-century America was the conflict between old-stock, primarily Protestant groups and more recently arrived, mostly Catholic immigrants. Particularly within the Democratic Party, these two competing groups found their respective champions in the reform wing and in the urban political machines. The machines offered immigrants food, medical attention, jobs, and assistance in navigating American life, sometimes ignoring legal niceties in the process; the corruption of the cities provided impulse for reform but also a convenient justification for reformers to indulge their ethnic, social class, and religious prejudices (Banfield and Wilson 1963, chaps. 9–11; Golway 2014; Hofstadter 1955, chap. 5).

By 1912 in New Jersey, Woodrow Wilson's strongest party support was concentrated among the reformers. Essex County's Democratic machine and its leader, Jim Smith, bitterly opposed Wilson, as we have seen. But populous Hudson County was the crucial prize. "As in all Democratic contests in New Jersey, Hudson County, the cornerstone of the state Democracy, was the center and pivot of the struggle" (Link 1947, 425). When Wilson ran for president the first time in 1912, Hudson's powerful Democratic Party was bitterly divided. The dominant progressive faction under the leadership of Mayor H. Otto Wittpenn backed Wilson (Boots 1917, 115). But a "regular" faction led by a young Frank

5. Atlantic and Cape May are the excluded Jersey Shore counties in their analysis, but the beach communities in those two counties were devastated by the lack of summer business and were far from relaxed about the sharks. Capuzzo (2001, 274) reported: "During the second week of July, the grand hotels, cottages, and guest house from Cape May north to Spring Lake reported an average of 75 percent vacancies on some of the best beach days of the year." That geographic range includes both Atlantic and Cape May counties. The member of Congress for Atlantic County introduced a bill appropriating \$5,000 to the federal Bureau of Fisheries to help in "the extermination of the man-eating sharks now infecting the waters of the Atlantic Ocean along the coast of New Jersey" (276). The town of Wildwood in Cape May County offered a \$1,000 bounty for the killing of a man-eating shark (Fernicola 2001, 81).

Hague worked to defeat him, spreading pamphlets across northern New Jersey cities with large numbers of immigrants “emphasizing Wilson’s alleged contempt for the ‘new’ immigrants from southern and eastern Europe,” especially Italian Americans (Link 1947, 425; see also Kelly 1961, 55, 112).

Once he became president, “Wilson distanced himself from direct involvement in New Jersey politics” (Greenberg 2012, 225). In consequence, by the spring of 1913 “Wilson and [his political lieutenant Joseph] Tumulty were losing command of New Jersey’s political destiny” (Blum 1951, 76). Most importantly, Hague took increasingly firm control of Hudson County in this period (Connors 1971, chap. 2; Hart 2013, chaps. 1–2). Because of Hudson’s dominance in Democratic votes, he quickly became the state’s party leader and a classic political “boss,” going on to rule for several decades as “the supreme dictator of the Democratic organization in the State” (Kerney 1926, 279; see also Connors 1971; Hart 2013; McKean 1940; Smith 1982).

Thus, Hague was much better placed in 1916 than he had been in 1912 to do Wilson harm, especially in those parts of northern New Jersey with substantial numbers of immigrants. Events strengthened his hand as well. World War I was underway, and although the United States had not yet joined the fight, Wilson’s lean toward the British side had become apparent. “The hostility of unregenerate German- and Irish-Americans to Wilson’s diplomatic and military policies confronted Democratic leaders at the Convention of 1916,” and Tumulty spent much of the campaign attempting to overcome that hostility (Blum 1951, 105, 105–9).⁶ Hughes, the Republican candidate, was endorsed by most major German and Irish American spokesmen and newspapers. Across much of the Northeast and upper Midwest, including New Jersey, urban Democratic machines catering to ethnic voters “either knifed the ticket or else they made only half-hearted campaigns” (Link 1954, 25).

What this review of the 1916 election in New Jersey demonstrates is that, compared to 1912, Wilson’s control of the Democratic Party was weaker, Hague’s was stronger, and anti-Wilson sentiments had grown in importance among many immigrant groups. These developments would be expected to depress Wilson’s vote among immigrants, particularly in places influenced by increasingly hostile Democratic political machines. Thus, in any analysis of how the vote changed

6. Irish Americans’ lack of enthusiasm for the prospect of America joining the war was compounded by an intensification of anti-British feeling after the Easter Rising in Dublin was put down and brutally punished in April 1916 (Link 1960, 20–23).

from 1912 to 1916, it is crucial to pay attention to the immigrant vote.⁷

Exactly how should one do that? In our book, we employed a dummy variable for Hague’s Hudson County plus the three counties adjacent to it, each of which had immigrant populations constituting more than 60% of its residents. One of those three counties is Essex, home of the other major Democratic machine in this period. The other two are Union and Bergen, parts of the same densely settled, immigrant-rich metropolitan area as the two well-known adjacent Democratic machine counties. We called these four the “machine counties.” However, it is certainly possible to imagine other historically “defensible,” “reasonable” approaches of the sort Fowler and Hall seek. Do those plausible alternatives support our conclusion about the effects of the shark attacks? The answer to that question turns out to be a firm yes.

We begin by supposing that a researcher knew nothing about political machines in New Jersey in this period but had a basic grasp of the issues in the 1912 and 1916 presidential elections and thus knew that immigrant voters were less happy with Wilson in 1916 than they had been in 1912. Hence, in a test for Wilson’s vote losses in the Shore counties, the proportion of immigrant voters in each county would need to be controlled. On its own, this is a somewhat naive specification, since it treats isolated immigrants in rural counties the same as those dependent on their local ward boss for employment.⁸ But it has the advantage of requiring only a standard census demographic variable, little political judgment, and no special knowledge of New Jersey. The result of that analysis is given in table 1, column 1.⁹ The estimated shark effect (the “beach” coefficient) is a bit larger than 3 percentage points, just as we argued in the book; the coefficient is statistically significant at the customary .05 level; and the overall goodness of fit is superior to both of Fowler and Hall’s proposed alternative specifications dealing with political machines.

A somewhat more sophisticated approach would take account of immigrant density. Where immigrants constituted less than 20% of the population in this period, as in rural counties like Ocean, Salem, or Warren, they were generally not organized by local machines, as the historical lit-

7. Happily, the 1910 US census provides county-level information about foreign-born citizens and about their children. Like most analysts, we counted anyone who was foreign-born or who had at least one foreign-born parent as an “immigrant.”

8. We found no effect of immigrant groups once the effect of the political machines was controlled (Achen and Bartels 2016, 124).

9. Votes and demographics are taken from the *Manual of the Legislature of New Jersey* and the US census, various years, just as in our book.

Table 1. The Beach Effect on Wilson's 1916 Vote

	(1)	(2)	(3)	(4)	(5)
Beach	-3.27** (1.50)	-3.50** (1.26)	-3.24** (1.27)	-3.41** (1.31)	-3.23*** (.99)
Urban		-3.90*** (1.14)	-6.11** (2.25)		
Immigrant proportion	-6.22* (3.04)		5.81 (5.12)		
Urban × immigrant				-5.36*** (1.72)	
Machine county					-5.66**** (1.09)
1912 Wilson vote (three-party share)	.885**** (.091)	.910**** (.076)	.937**** (.079)	.914**** (.079)	.947**** (.061)
Intercept	8.87* (4.55)	6.36* (3.50)	3.40 (4.34)	6.10 (3.63)	4.54 (2.78)
σ	2.48	2.11	2.09	2.20	1.70
Adjusted R^2	.864	.901	.903	.893	.936

Note. Dependent variable is the Democratic percentage of the two-party vote in New Jersey counties. Entries are regression coefficients (with standard errors in parentheses). $N = 20$ (Essex County excluded).

* Significant at .10.

** Significant at .05.

*** Significant at .01.

**** Significant at .001.

erature makes clear. (For the case of Ocean county, e.g., see McKean [1938], 183.) Thus, an appreciation for the political significance of immigrant mobilization would argue for distinguishing the counties where immigrants were heavily concentrated in this period (Vecoli 1965, 102): “over half of all persons of foreign birth in the state were to be found in the counties of Essex, Hudson, and Passaic; most of the rest were in Bergen, Middlesex, and Union counties. These were the only counties in 1900 in which the foreign stock accounted for more than 50 per cent of the total population. Native stock was in the majority in the other 15 counties. The six ‘foreign stock’ counties contained almost two-thirds of all the people in New Jersey, as well as more than three-quarters of the state’s foreign-born residents.”¹⁰

Alternatively, one might note that Democratic Party machines in this era generally flourished in urban areas. That consideration, even without regard to immigrant concentrations, would argue for paying special attention to the most densely settled counties. Thus, Reynolds’s (1988, 175) detailed study of New Jersey voting in this period distinguished six

counties with no rural areas from the rest of the Garden State counties, each of which still had rural areas in this period. Conveniently, the six urban counties identified by Reynolds are precisely the same six counties mentioned by Vecoli where immigrants were heavily concentrated—the four we have already identified as “machine counties” plus Passaic and Middlesex. Thus, an analyst cognizant of the significance of either ethnic tensions or urban political culture might well be led to expect a distinctive voting pattern in these six counties, even in the absence of any specific historical knowledge about the presence or activities of Democratic machines.

Focusing on these six urban, immigrant-rich counties provides only a rough proxy for machine influence in the 1916 election. Parts of Passaic and Middlesex are more than 40 miles from the Democratic strongholds of Newark and Jersey City, a considerable distance in this era. No party machine was organized in Middlesex until the late 1920s (McKean 1940, 56). Thus, our original judgment was that Passaic and Middlesex were in a different category from our four “machine” counties and that grouping all of them together did not get the politics of the period quite right. Nonetheless, for the moment we put aside the politics of the era and our own judgments; we simply adopt the definition of urban immigrant counties implied by census demographics and set out by other authors knowledgeable about this period of New Jersey history.

10. By 1910, immigrants made up more than 60% of the population in each of these six counties; no other county in the state had more than 50% immigrants.

In table 1, “urban” therefore refers to these six counties—Bergen, Essex, Hudson, Middlesex, Passaic, and Union. In column 2 of table 1 we report how the beach county effect holds up when “urban” is used as a control in place of our “machine” variable. The fit is improved from the naive specification in column 1 where only immigrant share was used, and again the beach county effect is above 3 percentage points and comfortably significant, this time with a p -value of .013.¹¹ Column 3 adds both the immigrant fraction and the urban variable to the regression, while column 4 uses their interaction. The beach coefficient is essentially unchanged in both cases, and again the effect is comfortably statistically significant.¹² Column 5 gives the original specification that we reported in our book. In our view, it takes the most careful account of the politics of the period; it also fits the data best. The beach county effect is very similar to that in all the other regressions, but it is now statistically significant at beyond the .01 level.¹³

The lag specifications in table 1 are not the only way to proceed. One could also compute the same regressions using the difference in Wilson’s vote share between 1916 and 1912 as the dependent variable. Then the naive regression paralleling column 1 produces a beach coefficient of -2.8 , significant at .10. The other regressions all show effects above 3 percentage points, and they are all statistically significant at the same levels as before.

One could also weight each observation in the regressions in table 1 by its 1916 vote total, giving more weight to the most populous counties.¹⁴ Weighting is probably not wise here for the reasons given by Deaton (1997, 66–73) and Solon, Haider, and Wooldridge (2015): the resulting statistical inefficiency is likely to outweigh the value of more appropriately averaging any parameter heterogeneity across counties. However, if one proceeds nonetheless, then all the estimated beach county effects are above 3 percentage points, and the number of asterisks for statistical significance is the same in each case as in the corresponding differenced regression.

Thus, across all of these regressions with various combinations of control variables, lag specifications, and weights,

11. Dropping either Passaic or Middlesex individually from the definition of “urban” makes no difference: the beach coefficient remains above 3 percentage points, with p -values less than .02 in both cases.

12. If immigrant, urban, and their interaction are all entered in the regression, collinearity makes it difficult to ascertain the effect of the three control variables, and the fit is no better. But the beach effect remains as before (-3.4), and the p -value is again less than .02.

13. If we add all the other control variables from table 1 to this regression, only the machine variable is statistically significant among them at .05, and the adjusted R^2 statistic indicates a slightly worse statistical fit; nonetheless, the beach county effect remains -3.2 with a p -value of .011.

14. Using the 1912 total votes as weights gives results that are virtually identical, as expected.

just one beach county parameter estimate is as low as -2.8 . The other estimates all lie between -3.0 and -3.7 . Two of the coefficients are significant at .10; the rest are significant at least at .05, and the most substantively plausible and best-fitting specifications show an effect significant at .01.¹⁵

In sum, the county data indicate a shark effect of about 3 percentage points, just as we said in our book. Unusually for aggregate data, a wide variety of sensible model specifications produce rather similar estimates of the beach coefficient. As Fowler and Hall’s own regressions demonstrate, even many historically inaccurate model specifications produce roughly similar results. The shark effect is simply very robust.

Fowler and Hall are concerned that “flexibility in data analysis can lead to false-positive results” (2018, XXX). That is a genuine problem, one that has been discussed in the statistical literature for 50 years or more. Examining how results “would have differed under other defensible specification choices” is a useful check on overconfidence (XXX). To that end, we discussed in our book an extensive set of alternate specifications and other statistical checks—and noted that our shark effect held up handsomely in all those well-fitting models. Fowler and Hall seem not to have noticed that evidence.

If Fowler and Hall’s alternative specifications were equally “defensible” substantively and fit equally well, the fact that they tend to produce slightly smaller and less precise estimates for the impact of the shark attacks would qualify our argument. But their statistical models fail on both counts. Generating plausible alternative specifications is hard work. In the current case, it requires substantive expertise about Progressive politics and attention to historical detail in the election returns. In our view, that is what good science requires.

Proliferating forking paths is easy. Substantively implausible, less well-fitting models that make the standard errors larger and the relevant coefficients statistically insignificant come readily to mind. Finding attractive, substantively plausible models is much harder. Some potential garden paths are known to lead to unstable footing or a collapsed bridge. Treating paths of that kind as attractive alternatives that should have been explored is both misleading and dangerous, and it does nothing to direct us to our destination.

15. We also computed Huber heteroskedasticity-robust standard errors and standard errors corrected for clustering, with the clusters defined variously as the Shore counties, the counties that border the Atlantic Ocean, or all the counties that border on water. All these raised the t -ratios and made the p -values for the beach effect even smaller. We do not take these asymptotic calculations to be exact values, but they do provide additional assurance about the consistent shark effect in table 1.

THE OCEAN COUNTY ANALYSIS

In the second part of our analysis, we looked at the vote in beach and near-beach townships in Ocean County. As we explained, we chose Ocean because the Jersey Shore townships there are located on two long thin strips of land separated from the mainland—the Barnegat Peninsula and Long Beach Island. In all cases, they are within five miles or so of mainland boroughs that are not Shore communities, and they are typically much closer. Apart from the beach communities having an Atlantic shoreline and summer visitors, the communities on each side are similar, with many citizens engaged in the same occupations (Anderson 1998; Hughes 1997; Stokley 2014; Wortman 1963). We showed that the beach and near-beach communities voted similarly in 1912 but quite differently in 1916, as the Shore punished Wilson for the lost income caused by the shark attacks.¹⁶ In particular, the two Shore towns that experienced shark attacks each reduced their support for Wilson by more than 8 percentage points, a remarkable drop that Fowler and Hall do not mention.

Fowler and Hall's criticism of our township-level analysis of election returns in Ocean County is of a rather different character from the rest of their critique. Although they raise a variety of questions about our data and analysis, the key issue here turns on the status of a single town, Seaside Park, which we omitted from the analysis on historical grounds. Adding it back into our analysis, as Fowler and Hall do, reduces the apparent effect of the shark attacks on Wilson's vote share by more than half and the goodness of fit (R^2) by almost four-fifths. However, their argument for including Seaside Park in the analysis rests on a historical misjudgment. By carefully examining the history and politics of

Ocean County, we corroborate our original treatment of this case and the resulting conclusion that the townships most directly affected by the shark attacks turned sharply against Wilson at the polls.

We begin by noting that much of the Ocean County shore consisted of empty sand dunes at the turn of the twentieth century, as many local histories remark (e.g., Wilson 1953, chap. 21). But as the railroads began to reach Shore areas, more residents arrived. Sometimes borough or township boundaries were adjusted as real estate companies bought unpopulated tracts of land for development. Fowler and Hall worry about these boundary adjustments affecting vote totals, as they might if they were done in recent decades, but in the period under study, these were not instances of moving substantial numbers of people between voting units. Typically the land being reallocated was nearly or entirely empty.¹⁷

The exception to this pattern occurred in the unincorporated area north of Seaside Park, which eventually became the borough of Seaside Heights. In 1909, Seaside Heights was "a barren tract owned by the Manhasset Realty Company" (Wortman 1963, 24, 52). The owners were prosperous Camden-area businessmen with plans for a major resort, not just a few home sales. In 1913, the New Jersey legislature was induced to permit the area to become its own borough and separate from Berkeley Township, an area almost entirely located on the mainland but of which Seaside Heights was a legal part.¹⁸ Thus, Seaside Heights has vote returns reported for the second Wilson election in 1916 but not for the first in 1912, when it did not yet exist.

The most pertinent argument that Fowler and Hall make in their discussion of our Ocean County results concerns this new borough of Seaside Heights. We stated in our book (Achen and Bartels 2016, 126 n. 16) that the older, adjacent borough of Seaside Park "apparently split into two between 1912 and 1916 and jointly nearly doubled in size" as Seaside Heights was formed. Hence, the 1916 voting population in Seaside Park was not at all comparable to the 1912 electorate before the split, and we deleted Seaside Park from our analysis for that reason, just as we did other townships with more than 25% growth in that era of rapid Jersey Shore development.

16. Fowler and Hall attempt to replicate this comparison along the rest of the Jersey Shore, but their details are skimpy. To ensure meaningful comparisons, did they drop areas whose populations grew substantially? Similarly, were they able to closely match beach and nonbeach townships and boroughs on prior voting behavior? Did they have the same difficulties distinguishing classic Jersey Shore communities from non-Shore areas that they report having in Ocean County? What did they do about the Garden State Parkway and its sharp turn away from the beach in Monmouth County, making it a much poorer dividing line between beach and nonbeach areas there than it is elsewhere? How did they code Neptune Township, which connects to the Shore only at Ocean Grove, a teetotaling Methodist campground that effectively discouraged all summer visitors except those who came to hear sermons and to pray? And so on. As we try to demonstrate in the following paragraphs, details of that kind really matter. Making this sort of comparison persuasively in the book for just one county (Ocean), an unusually tractable case, required many weeks of patient historical investigation on our part, including travel to the area to consult the county library and official records, as well as visits to the State Archives in Trenton. In the absence of details about their research decisions, we are reluctant to hazard an assessment of Fowler and Hall's thinly reported work on three other, much less tractable counties.

17. That is the case for Beach Haven, e.g., which Fowler and Hall drop or merge with the adjacent town of Long Beach in some of their analyses because of concerns about boundary shifts (*Ocean County Courier* 1913f).

18. This new borough was authorized by the New Jersey legislature in February 1913, subject to a referendum of the inhabitants (Legislature of the State of New Jersey 1913). That referendum was held on March 25, 1913, and township status was approved 21–0 (*Ocean County Courier* 1913g). In April, township officers were elected, and the township government came into being (*Ocean County Courier* 1913a; Wortman 1963, 25).

Fowler and Hall note correctly, however, that in the period before 1913, unincorporated Seaside Heights was not legally part of Seaside Park. Instead, it belonged to Berkeley Township. This leads Fowler and Hall to assume that Seaside Heights votes were counted separately from Seaside Park, and that they were included in the Berkeley Township totals before 1913. In that case, Seaside Park votes in 1912 and 1916 came from the same borough, with no Seaside Heights votes included at either election. And since Seaside Park alone did not have more than a 25% increase in votes cast, on this argument it should have been included in our analysis. The result, as Fowler and Hall show, would be to dramatically reduce the apparent impact of the shark attacks in the beach townships of Ocean County.

From the perspective of today's procedures for election administration, Fowler and Hall's assumption that Seaside Heights voted in Berkeley Township would make considerable sense. As it happens, a similar unincorporated area, South Seaside Park, is adjacent to Seaside Park even today. It, too, belongs to Berkeley Township, and its citizens' votes are included there. Thus in recent decades, being an unincorporated area on the Barnegat Peninsula that belongs to Berkeley Township has meant just what Fowler and Hall imagine that it meant in 1912. However, just as statisticians by occupation worry about overfitting, historians by occupation worry about "presentism," the danger of assuming that the past was just like the present. The issue about Seaside Park, then, is this: Can we assume that the present municipal procedures in Berkeley Township tell us how things were done a century ago?

To answer that question, it is helpful to remind oneself what life was like on the Barnegat Peninsula in the early twentieth century. The first horse and auto bridge to the Seaside Park area, a wooden structure, was not opened until October 1914, so that getting from isolated communities like Seaside Park and Seaside Heights to the mainland in earlier years was no simple matter. A few people had automobiles, as the *Ocean County Courier* newspaper frequently noted at the time. But most visitors and residents relied on trains: "In the first two decades of the century, however, Shore residents continued to rely primarily on the railroad for transport" (Wilson 1953, 847; also Wortman 1963, 9).

In 1912, the sole Berkeley Township polling place was in Bayville borough, as the official polling place announcements in the newspapers show (*Ocean County Courier* 1912) and as the New Jersey State Archives confirm. There was no precinct in unincorporated Seaside Heights. The total number of precincts in Ocean County in 1912 was 32, and they are listed in the *Manual of the Legislature of New Jersey* (Fitzgerald 1913, 126–27 and 644). Seaside Heights is not among them. No votes were recorded there, as we confirmed from the

handwritten returns from each polling place at the New Jersey State Archives.¹⁹ When elections were held in Seaside Heights the following year, the number of precincts had increased by one, and Seaside Heights appears on the list (Fitzgerald 1914, 594, 616). It was noted in 1913 that those were the first elections ever conducted in that borough (*Ocean County Courier* 1913a).

Thus, if Seaside Heights voters were included in Berkeley Township, as Fowler and Hall suppose, their votes would have been cast in Bayville. But reaching Bayville from Seaside Heights would have required two separate railroad connections covering a distance of more than 11 miles (Discover Seaside Heights 2017) and then a wait for the connecting trips back to Seaside Heights. The whole expedition would have cost substantial time and money.

Alternatively, one could take on the expense of hiring a boat to cross Barnegat Bay (Miller 2000, 509), but that would still have left voters several miles from Bayville, and they would have needed to hire a horse (or have one towed across the bay with them on a barge) or face a long walk (Stokley 2014, 49). The only other alternative would be to travel by horse or auto up the peninsula to Bayhead to reach the mainland and then back down to Bayville, a distance of 30 miles (Ocean County Principals' Council 1940, 151). Any of these journeys would have required a dramatic effort on election day, which was not inexpensive, consuming hours in each direction. Neither the *Asbury Park Press* nor the *Ocean County Courier* newspaper at the time ever mentions a dramatic he-gira of this sort when Seaside Park and adjacent area election events are reported, even though much shorter and easier trips from South Seaside Park across the new bridge in the post-World War II era receive frequent coverage for their inconvenience to the citizenry.²⁰

19. A part of what became Seaside Heights belonged to Dover Township, but Berkeley Township is the relevant unit for electoral purposes. When Seaside Heights was incorporated, the legislature designated Berkeley as the jurisdiction to manage the transition, and the election notice was issued by the Clerk of Berkeley Township (*Ocean County Courier* 1913c). In any case, Dover Township had just three polling places in November 1912, all in the town of Toms River (*Ocean County Courier* 1912). The original handwritten records in the State Archives show that no votes from Seaside Heights were reported from any Toms River precinct.

20. In 1975, a majority of South Seaside Park's residents sued the Berkeley Township Committee to secede from Berkeley and join Seaside Park. Their children were being bused to schools in Bayville borough in Berkeley Township. Even with the modern bridge erected in 1950, the trip took the children an hour each way. Citizens dealing with property taxes or other business had to travel to the township offices in Bayville, too (Colford 1975). South Seaside Park had developed a modest population only in the 1920s, after the first bridge was built in 1914, and thus the township could insist that they come to Bayville for township business. The township enforced the rule that their votes

So if the citizens of Seaside Heights did not undertake the elaborate trek to Berkeley Township in the years before the construction of the bridge and did not vote in their own unincorporated area, where did they cast their votes?

The only remaining option is that they joined their neighbors next door in Seaside Park. Accounts of this period demonstrate that township boundaries were not enforced as rigidly as they are now. In the period before incorporation, Seaside Heights was often referred to as North Seaside Park or Seaside Park North—just the northern part of Seaside Park (*Ocean County Courier* 1913d). The initial legal notice of planned appeal to the legislature for incorporation of Seaside Heights noted that it would include the area “known as the plan of North Seaside Park,” and the official act of the legislature authorizing the new borough used the same phrase (Legislature of the State of New Jersey 1913; *Ocean County Courier* 1913b). This verbal description corresponded to daily life for the few dozen inhabitants. Residents of that area got their mail in Seaside Park, for instance, and did not get their own post office until after incorporation (*Ocean County Courier* 1913e; Ocean County Principals’ Council 1940, 151). They shopped and attended church services in Seaside Park as well, since businesses developed only slowly in Seaside Heights, and no churches were established until 1913 (Ocean County Principals’ Council 1940, 151; Wortman 1963, 24, 50). Moreover, Seaside Heights children went to school in Seaside Park until their own school was built in 1913, even though they lived in an entirely different township (Borough of Seaside Park 2016; Ocean County Principals’ Council 1940, 151). With no bridge in place, Berkeley Township was simply too distant. Thus, even in a relatively expensive matter such as school enrollment, in this time of very small towns with limited communication and close personal ties, living a few blocks away and being an integral part of borough life was thought to be a more relevant consideration than formal township boundaries.

Perhaps most persuasively, the act of the legislature establishing procedures for the first vote in Seaside Heights explicitly instructed Berkeley Township to create a voter roll for the incorporation election. The township had such a list for its own voters (the 1912 presidential canvass having been conducted just a few months before) but did not have one for Seaside Heights (Legislature of the State of New Jersey 1913). Obviously, then, before 1913 Seaside Heights residents were not voting in Berkeley. Contrary to what Fowler and Hall

had to be cast in Bayville, too, until they got their own polling place in 1946 (*Asbury Park Press* 1946a, 1946b). Discussion of this problem in the local newspapers from the 1940s onward is extensive, and secession efforts continue down to the present day (Miller 2015).

suppose, their votes had been cast in the adjacent borough of Seaside Park, just as they got their mail, sent their children to school, and handled their other personal business there. Thus, the historical evidence forcefully supports the view that votes from Seaside Park and what became Seaside Heights were counted jointly in 1912 and reported as “Seaside Park,” as we argued in our book.

Fowler and Hall proceed by ignoring or overriding this historical evidence. But even without attention to the history, careful consideration of the quantitative evidence would have provided ample warning not to proceed as Fowler and Hall do. The Philadelphia businessmen who founded Seaside Heights were coming from a city with a GOP machine, in a period when business elites favored Republicans in any case. Unsurprisingly, then, when the 1913 New Jersey gubernatorial race took place, the first election held in the borough, the Seaside Heights vote tilted sharply Republican. Edward Stokes, the GOP candidate, got just under 60% of the three-party vote, and the Democrat only 30%.

Now, in Ocean County as a whole, 1913 was a rerun of 1912: Wilson’s share in 1912 and the share of his handpicked Democratic candidate, James Fielder, in 1913 were virtually identical (38.5% vs. 38.8%). Yet the Democratic Party share in Seaside Park took a dramatic 10 percentage point jump upward in 1913.²¹ If we examine the four Wilson-related elections in this period—Wilson’s gubernatorial race in 1910, his two presidential elections, and the gubernatorial contest in 1913 with his chosen successor, Fielder (Blum 1951, 76–77), the sharp break between 1912 and 1913 in Seaside Park and Seaside Heights becomes obvious, as figure 1 demonstrates.²² What happened?

The obvious interpretation, of course, is that heavily GOP Seaside Heights had been part of Seaside Park in 1912, just as the historical evidence suggests. Removing that mass

21. The only other geographic unit in our data set of any size that moved substantially toward Wilson (although less than Seaside Park) in this period is Ocean Township, which is not on the Shore. It was trending steadily Democratic throughout 1911–16, probably because of demographic changes. Dropping this borough from our original analysis makes no difference; the resulting shark effect on the Ocean County Shore becomes 9 percentage points, statistically significant at .05.

22. Paralleling the approach in our county analysis, the entries in the figure are the two-party Democratic shares in 1910 and 1916, along with the three-party shares in 1912 and 1913 when the Progressives received substantial votes in Ocean County. If instead of the gubernatorial contest in 1913, we had used the arguably less relevant Senate race (where Wilson also intervened on behalf of the winning progressive primary candidate, William Hughes, but was not involved in his initial selection, according to Link [1947, 497–98]), the Democratic fraction in each borough would rise by a little over 9 percentage points, maintaining the same size break between the two boroughs seen in fig. 1.




Figure 1. Woodrow Wilson’s vote shares before and after the 1913 incorporation of Seaside Heights.

of Republican voters left a greater proportion of Democrats in Seaside Park, and so its Democratic fraction rose in 1913. Both communities then maintained their new partisan balance in 1916, as figure 1 shows. This is all just as one would expect if the two communities’ votes were counted jointly at the previous election—but wholly mysterious on Fowler and Hall’s view that Seaside Park was unchanged throughout this period.

Turnout figures tell much the same story. The Barnegat Peninsula, where both Seaside Park and Seaside Heights are located, was gaining population steadily in this period. Seaside Park had 54 voters in 1910 and 90 in 1916; Seaside Heights went from 27 upon incorporation in 1913 to 69 in 1916. Yet the vote cast in Seaside Park, after rising steadily in prior years, suddenly dropped in the high-salience gubernatorial election of 1913 below the level recorded in a much less visible election for the state assembly in 1911, and it remained below that level for several years. As figure 2 shows, the 1913 drop in Seaside Park vote totals occurred just when Seaside Heights was incorporated—precisely what we would expect if votes from Seaside Heights had formerly been counted in Seaside Park.

Meanwhile, Berkeley Township, which Fowler and Hall suppose had lost votes to Seaside Heights, does just the opposite. Unlike Seaside Park, its turnout jumped by 18% between 1911 and 1913, as one would expect when a geographic unit is unchanged and a low-visibility election is succeeded by a more salient gubernatorial election.²³

23. In fig. 2 and in comparisons to 1911, we use the largest total two-party vote for any reported office as our measure of turnout (except for 1912 and 1913, with strong Progressive candidates running, where we use the three-party total). In figs. 2 and 3 for 1914, we use the congressional vote because it is the highest office total in Seaside Park and in the combined vote of the two Seasides. We ignored the fact that in Seaside Heights there are four more votes for state senator and one more for Assembly in 1916 than for president. Obviously, none of these tiny differences makes any difference in our conclusions.

Finally, figure 3 shows the combined turnout in Seaside Park and Seaside Heights from 1910 through 1916. By contrast to the erratic path of total turnout when Seaside Heights is pulled out from Seaside Park, if we instead combine the two communities the trend line of total votes cast traces a smooth upward path, just as one would expect from the steady population growth in this period.

In sum, even if one focuses solely on the quantitative evidence, scarlet billows start to spread around Fowler and Hall’s interpretation. This brings us back to our original judgment: Seaside Park’s reported vote totals in 1912 and 1916 are not comparable, and thus it should be dropped from the statistical analysis. That returns us to the conclusion we reported in our book, with all the specification checks we set out there and the same drop in Wilson’s vote in the Ocean County shore communities of about 10 percentage points.

PLACEBO TESTS

Finally, Fowler and Hall produce a series of “placebo tests” intended to test whether “beach and nonbeach counties could be subject to different idiosyncratic shocks, making differences likely even in the absence of shark attacks” (2018, XXX). Since New Jersey alone does not provide enough cases for this sort of analysis, they compare electoral swings in coastal and noncoastal counties in every state with ocean shorelines in every presidential election since 1872. They find that 27% of these differences are statistically significant, so that a negative effect in the Shore counties, as we found, would occur about one time in seven.

Fowler and Hall acknowledge that “we cannot exactly replicate Achen and Bartels’s strategy in every state election” (2018, XXX). Indeed, expanding the analysis from a single election in a single state to three dozen elections in each of 20 states raises a variety of daunting complications. As a simple mechanical matter, the probability of observing “significant”




Figure 2. Total votes cast in Seaside Park and Seaside Heights, 1910–16




Figure 3. Combined turnout by year for Seaside Park and Seaside Heights, 1910–16.

differences in election outcomes between coastal and noncoastal counties in a given state depends crucially on how many coastal and how many noncoastal counties there are; Fowler and Hall make no attempt to adjust for differences of that sort. More substantively, the characteristics of coastal and noncoastal counties differ greatly from state to state, and they make no attempt to account for those differences either. Nor do they take any account of political factors other than coastal status that might matter in a given election in a given state. It is hardly surprising that this style of apolitical analysis will sometimes produce “false positive” results. But that recognition should be the beginning, not the end, of serious data analysis. In particular, it should be the beginning of a search for credible alternative explanations.

Omitted factors are always a concern in analyses of the sort we presented. As we noted in the book, we spent a good deal of time considering potential confounding factors familiar from the literature of the time and from subsequent historical studies. None of them eliminated the shark effect. Similarly, having verified that coastal counties in other times and places have sometimes behaved distinctively at the polls, talented political scientists like Fowler and Hall might have asked themselves where and when such differences have occurred, what are the many factors that might have caused them, and—most importantly for current purposes—what implications, if any, those alternative political factors might have for the Jersey Shore in 1916. Were one or more of those factors operating powerfully in the Shore that year but not elsewhere in New Jersey? Does including them in the analysis eliminate the shark effect? In other words, how does the politics work in an alternative explanation?

Fowler and Hall do not pursue any part of that task. They simply point to the occasional effects in other years and other states of their generic “placebo,” a grab bag of all conceivable

causes of electoral shifts, as grounds for doubting our finding. But noticing the familiar fact that coastal counties sometimes vote distinctively in other states and other elections falls far short of providing a sensible counter to our claim of a shark effect on the Jersey Shore in 1916. It is mere “hit and run” criticism (Bross 1960)—a vague hypothesis that “maybe something else happened.” But that is not enough for credibility.

For authors and critics alike, the task is always to provide a credible substantive argument and show that it empirically dominates plausible alternatives. That is what we tried to do in our book. By that standard, the evidence forcefully favors our conclusion that the shark attacks mattered.²⁴

CONCLUSION

Broad concerns of the sort raised by Fowler and Hall’s critique of our work are always in order. Are the data pertinent and the analyses replicable? Do the findings mean what the analysts say they mean? Are the assumptions and judgments underlying the statistical work substantively plausible? Do the results hold up under other substantively plausible approaches and interpretations? Or were the analyses consciously or unconsciously constructed to support a particular argument? Human beings are always subject to errors and biases, and good science requires both open-mindedness and humility.

Yet critics raising those questions are not always persuasive in answering them, especially when their criticisms are grounded in generic methodological skepticism rather than careful engagement with the relevant theory and evidence. We have tried to show here why we adhere to our original finding that the 1916 New Jersey shark attacks reduced Woodrow Wilson’s presidential vote in the economically affected areas that fall. Both the county data and the township data from Ocean County support that finding. And we have shown, both

24. Fowler and Hall (2018, fig. 3) also supply a graph meant to show that 1912 is an outlier among presidential elections in the Progressive period, in that the Democratic vote share shows a wider gap between the shore counties and the rest of the state in that year than in the others. This might suggest that the relevant event happened in 1912, not in 1916 when the shark attacks took place. Fowler and Hall do not pursue the substantive question of what that 1912 event might be. Moreover, to anyone familiar with New Jersey voting in this period, their graph will look odd in other respects; for example, New Jersey’s steady drift toward the Democrats as immigrants pour in after World War I does not appear in the graph. In fact, we can reproduce their graph only by averaging the vote shares across counties without weighting by population. In this era of rapid population growth in some counties, and huge differences in population across counties, one always has to check whether a result is robust to weighting by population. In this case, after population weighting, 1912 and 1916 look very similar, and 1912 ceases to be an outlier, which eliminates Fowler and Hall’s supposed anomaly. The steady trend of the nonbeach counties to the Democrats is also visible in the graph once the votes are properly weighted.

here and in our book, that the finding is robust against the important plausible alternative approaches and interpretations of the evidence.

Some readers may be surprised by the amount of detail needed to make our case. Can we not just get on with running regressions and forget about all that historical stuff? But the main task of historical research is always to get right the culture and politics of the time and place being studied, and that requires diligent effort to overcome one's erroneous preconceptions. L. P. Hartley's novel, *The Go-Between*, has as its well-known first line, "The past is a foreign country: they do things differently there." Taking the implicit admonition seriously is what we tried to do in our study of New Jersey in the Wilson era. Getting the past right is not quick work, and like everyone else who does this kind of research, we have learned that the methodological frameworks that predominate in statistics, economics, public policy, and business departments, while profoundly helpful, are not in themselves sufficient to prevent inferential blunders in historical studies. Thus, we did our best to get the history right and to ground our argument in the full range of available evidence, both quantitative and qualitative.

Like all scientific work, of course, historical analysis is always provisional pending additional evidence and insight. In the case of the shark attacks, New Jersey voting returns at the borough level are available for the entire state. Some census data are available for each borough; in larger cities, both electoral and some demographic data are disaggregated by wards. Moreover, for many cities and counties of the state, there are local histories, volumes of political reminiscences, and books of recollections by residents, only a few of which we have consulted for counties other than Ocean. Both Wilson's and Joseph Tumulty's papers are available to scholars. Although many of these sources are available neither on the Internet nor in university holdings, they can be found in city and county libraries or, in the case of Wilson's papers, in federal archives. Serious future scholarship may well exploit some or all these sources to elaborate or modify our findings. But there is no convenient shortcut to more reliable inferences.

In recent years, quantitative political scientists have become concerned, some would say obsessed, with the inferential risks of nonexperimental data. (The inferential risks of experimental data—particularly the daunting challenge of demonstrating external validity—have attracted much less attention lately, although they, too, have been well known for decades.) Learning how to make better, more reliable inferences is always a worthy goal, and we have devoted substantial fractions of our careers to that pursuit. But the way forward is not to be found solely within the inferential templates provided by statistics texts.

Fowler and Hall set out to demonstrate "that recent concerns about voter competence, at least as they relate to voters' abilities in ignoring irrelevant events, are overblown" (2018, XXX). They argued that we could have done our analysis many other ways and that some of those other "forking paths" would have brought us out differently. But simply multiplying forking paths without regard to their theoretical relevance or substantive plausibility is a distinctly unreliable approach precisely because it is so easy. In their final sentence, Fowler and Hall cast their criticism of our work as a parallel to Fowler and Montagnes's (2015) critique of Healy, Malhotra, and Mo's (2010) analysis of the electoral impact of college football games. In our view the parallel undercuts their argument rather than bolstering it. While Fowler and Montagnes argue that the Healy et al. finding "is most likely a false positive" (Fowler and Hall 2018, XXX), that conclusion has been repeatedly punctured by analyses employing a variety of different research designs and data tending to confirm the original finding (Busby, Druckman, and Fredendall 2017; Healy, Malhotra, and Mo 2015; Miller 2013). Fowler and Hall do not pause to consider that evidence, but for anyone who does, the parallel is more likely to raise concern about the potential for false negative inferences stemming from substantively meaningless multiplication of forking paths rather than about the potential for false positive inferences in our work.

Perhaps the saddest instance of how statistical shibboleths unmodified by substantive knowledge can go wrong is provided by the great statistician R. A. Fisher (1957, 1958a, 1958b), who argued that smoking probably did not cause lung cancer. Fisher had pioneered randomized trials, and he was unable to think past them. With no academic training in oncology, he pontificated on the subject as an autodidact. Even in Fisher's own time, a better-informed biostatistician, working with several prominent oncologists and epidemiologists, was able to show that the evidence strongly favored a carcinogenic effect for tobacco smoke (Cornfield et al. 1959). But Fisher persisted. The tobacco companies were happy to support his work as they battled the proposed laws and regulations whose enactment would ultimately save millions of lives (Bodmer 2003, 941).²⁵ Naive about human judgment, not least his own, Fisher worried that inferring the causes of cancer from observational evidence violated his statistical norms. He wrote, "I do not relish the prospect of this science being discredited by a catastrophic and conspicuous howler" (Fisher 1957, 298). But the catastrophic error—both professional and moral—turned out to be his.

25. Fisher was by no means alone among statisticians and econometricians with no expertise in oncology who took money from tobacco companies during the second half of the twentieth century.

As the wisest statisticians have always recognized, persuasive empirical science does not come from applying abstract statistical considerations to poorly grasped research problems. Rather, it emerges from deep substantive knowledge in dialogue with relevant statistical theory. That was the standard to which we aspired in our book. In our view, that kind of thinking, imperfect and provisional as always, represents the way forward for empirical political science.

ACKNOWLEDGMENTS

We offer special thanks to those who provided invaluable help at various stages of our research on early twentieth-century New Jersey elections, including the Ocean County Office of the County Clerk; the staff at the Ocean County Library, especially Colleen Goode; and the staff at the New Jersey State Archives, especially Bette Epstein. We are also grateful to Rogers Smith and to a thoughtful anonymous reviewer for their suggestions.

REFERENCES

- Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton, NJ: Princeton University Press.
- Anderson, Andrew J. 1998. *Seaside Park*. Charleston, SC: Arcadia.
- Apter, David E. 1964. "Ideology and Discontent." In David E. Apter, ed. *Ideology and Discontent*. New York: Free Press, 15–46.
- Asbury Park Press. 1946a. "26,200 Ocean Voters Listed." *Asbury Park Press*, October 17.
- Asbury Park Press. 1946b. "Berkeley Allowed 3d Voting District." *Asbury Park Press*, June 21.
- Banfield, Edward C., and James Q. Wilson. 1963. *City Politics*. New York: Vintage.
- Bartels, Larry M. 1997. "Specification Uncertainty and Model Averaging." *American Journal of Political Science* 41 (2): 641–74.
- Berg, A. Scott. 2013. *Wilson*. New York: Putnam.
- Blum, John M. 1951. *Joe Tumulty and the Wilson Era*. Boston: Houghton Mifflin.
- Bodet, Marc Andre, Melanee Thomas, and Charles Tessier. 2016. "Come Hell or High Water: An Investigation into the Effects of a Natural Disaster on a Local Election." *Electoral Studies* 43 (September): 85–94.
- Bodmer, Walter. 2003. "R. A. Fisher, Statistician and Geneticist Extraordinary." *International Journal of Epidemiology* 32 (6): 938–42.
- Boots, Ralph Simpson. 1917. *The Direct Primary in New Jersey*. New York: New Jersey Bureau of State Research.
- Borough of Seaside Park. 2016. "History." Borough of Seaside Park, New Jersey. <https://www.seasideparknj.org/history/> (accessed December 11, 2016).
- Bross, Irwin D. J. 1960. "Statistical Criticism." *Cancer* 13 (2): 394–400.
- Busby, Ethan C., James N. Druckman, and Alexandria Fredendall. 2017. "The Political Relevance of Irrelevant Events." *Journal of Politics* 79 (1): 346–50.
- Cantril, Hadley. 1941. *The Psychology of Social Movements*. New York: Wiley.
- Capuzzo, Michael. 2001. *Close to Shore: A True Story of Terror in an Age of Innocence*. New York: Broadway.
- Colford, Paul D. 1975. "South Seaside Park Is Suing for Secession from Berkeley." *New York Times*, December 7.
- Connors, Richard J. 1971. *A Cycle of Power: The Career of Jersey City Mayor Frank Hague*. Metuchen, NJ: Scarecrow.
- Cooper, John Milton, Jr. 2009. *Woodrow Wilson*. New York: Vintage.
- Cornfield, Jerome, William Haenszel, E. Cuyler Hammon, Abraham M. Lilienfeld, Michael B. Shimkin, and Ernst L. Wynder. 1959. "Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions." *Journal of the National Cancer Institute* 22 (1): 173–203.
- Deaton, Angus. 1997. *The Analysis of Household Surveys*. Baltimore: Johns Hopkins University Press.
- Discover Seaside Heights. 2017. "All Aboard: The Importance of the Railroads." Discover Seaside Heights (New Jersey). <http://www.discoverseasideheights.com/history/all-aboard>.
- Edelman, Murray. 1964. *The Symbolic Uses of Politics*. Urbana: University of Illinois Press.
- Edelstein, Jeff. 1999. *The Best of Everything at the Jersey Shore*. Morristown: New Jersey Monthly.
- Fernicola, Richard G. 2001. *Twelve Days of Terror: A Definitive Investigation of the 1916 New Jersey Shark Attacks*. Guilford, CT: Lyons.
- Fisher, Ronald A. 1957. "Dangers of Cigarette-Smoking." *British Medical Journal* 2 (5039): 297–98.
- Fisher, Ronald A. 1958a. "Cancer and Smoking." *Nature* 182:596.
- Fisher, Ronald A. 1958b. "Cigarettes, Cancer, and Statistics." *Centennial Review* 2:151–66.
- Fitzgerald, Thomas F. 1886–1950. *Manual of the Legislature of New Jersey*. Trenton: Legislature of the State of New Jersey.
- Fowler, Anthony, and Andrew B. Hall. 2018. "Do Shark Attacks Influence Presidential Elections? Reassessing a Prominent Finding on Voter Competence." *Journal of Politics* 80 (4): XXX–XXX.
- Fowler, Anthony, and B. Pablo Montagnes. 2015. "College Football, Elections, and False-Positive Results in Observational Research." *Proceedings of the National Academy of Sciences of the USA* 112 (45): 13800–4.
- Gaventa, John. 1980. *Power and Powerlessness*. Urbana: University of Illinois Press.
- Golway, Terry. 2014. *Machine Made: Tammany Hall and the Creation of Modern American Politics*. New York: Liveright.
- Gosnell, Harold F., and Norman Pearson. 1939. "The Study of Voting Behavior by Correlational Techniques." *American Sociological Review* 4 (6): 809–15.
- Gould, Lewis L. 2008. *Four Hats in the Ring*. Lawrence: University Press of Kansas.
- Gould, Lewis L. 2016. *The First Modern Clash over Federal Power: Wilson vs. Hughes in the Presidential Election of 1916*. Lawrence: University Press of Kansas.
- Greenberg, Brian. 2012. "The Progressive Era." In Maxine N. Lurie and Richard Veit, eds., *New Jersey: A History of the Garden State*. New Brunswick, NJ: Rutgers University Press, 202–35.
- Hart, Steven. 2013. *American Dictators: Frank Hague, Nucky Johnson, and the Perfection of the Urban Political Machine*. New Brunswick, NJ: Rutgers University Press.
- Healy, Andrew J., Neil Malhotra, and Cecilia Hyunjung Mo. 2010. "Irrelevant Events Affect Voters' Evaluations of Government Performance." *Proceedings of the National Academy of Sciences of the USA* 107 (29): 12804–9.
- Healy, Andrew, Neil Malhotra, and Cecilia Hyunjung Mo. 2015. "Determining False-Positives Requires Considering the Totality of Evidence." *Proceedings of the National Academy of Sciences of the USA* 112 (48): E6591.
- Hofstadter, Richard. 1955. *The Age of Reform: From Bryan to F.D.R.* New York: Vintage.
- Hosford, Hester E. 1912. *Woodrow Wilson and New Jersey Made Over*. New York: Knickerbocker.
- Hughes, Kevin. 1997. *Barnegat*. Charleston, SC: Arcadia.
- Kelly, Frank K. 1961. *The Fight for the White House: The Story of 1912*. New York: Crowell.

- Kerney, James. 1926. *The Political Education of Woodrow Wilson*. New York: Century.
- Kramer, Gerald H. 1971. "Short-Term Fluctuations in U.S. Voting Behavior, 1896–1964." *American Political Science Review* 65 (1): 131–43.
- Legislature of the State of New Jersey. 1913. Chapter 29 in *Acts of the One Hundred and Thirty-Seventh Legislature*. Trenton: Legislature of the State of New Jersey.
- Link, Arthur S. 1947. *Wilson: The Road to the White House*. Princeton, NJ: Princeton University Press.
- Link, Arthur S. 1954. *Woodrow Wilson and the Progressive Era, 1910–1917*. New York: Harper.
- Link, Arthur S. 1960. *Wilson: The Struggle for Neutrality, 1914–1915*. Princeton, NJ: Princeton University Press.
- Link, Arthur S. 1965. *Wilson: Campaigns for Progressivism and Peace, 1916–1917*. Princeton, NJ: Princeton University Press.
- Lovell, S. D. 1980. *The Presidential Election of 1916*. Carbondale: Southern Illinois University Press.
- Ludington, Arthur. 1911. "Election Laws: The New Geran Law in New Jersey." *American Political Science Review* 5 (4): 579–85.
- Mayhew, David R. 1986. *Placing Parties in American Politics: Organization, Electoral Settings, and Government Activity in the Twentieth Century*. Princeton, NJ: Princeton University Press.
- McKean, David Dayton. 1938. *Pressures on the Legislature of New Jersey*. New York: Columbia University Press.
- McKean, David Dayton. 1940. *Boss: The Hague Machine in Action*. Boston: Houghton Mifflin.
- Miller, Michael K. 2013. "For the Win! The Effect of Professional Sports Records on Mayoral Elections." *Social Science Quarterly* 94 (1): 59–78.
- Miller, Patricia A. 2015. "The First of Many." *Berkeley Patch*, January 9. <https://patch.com/new-jersey/berkeley-nj/south-seaside-park-secession-hearing-only-first-many> (accessed January 20, 2017).
- Miller, Pauline S. 2000. *Ocean County: Four Centuries in the Making*. Toms River, NJ: Ocean County Cultural and Heritage Commission.
- Nathan, Meyer Jonah. 1965. "The Presidential Election of 1916 in the Middle West." PhD diss., Princeton University.
- Ocean County Courier*. 1912. "Notice of Elections, Ocean County, New Jersey." *Ocean County Courier*, November 1.
- Ocean County Courier*. 1913a. "First Election at Seaside Heights." *Ocean County Courier*, April 25.
- Ocean County Courier*. 1913b. "Notice." *Ocean County Courier*, January 17.
- Ocean County Courier*. 1913c. "Notice of Election under Law Incorporating Borough of Seaside Heights." *Ocean County Courier*, March 7.
- Ocean County Courier*. 1913d. "Preparing Place for Bridge to Beach at Seaside Park North." *Ocean County Courier*, January 24.
- Ocean County Courier*. 1913e. "Seaside Heights." *Ocean County Courier*, May 2.
- Ocean County Courier*. 1913f. "To Enlarge Beach Haven Borough." *Ocean County Courier*, February 7.
- Ocean County Courier*. 1913g. "Unanimous for Borough at Seaside Heights Election." *Ocean County Courier*, March 28.
- Ocean County Principals' Council. 1940. *Tides of Time in Ocean County*. Philadelphia: Howe.
- Reynolds, John F. 1988. *Testing Democracy: Electoral Behavior and Progressive Reform in New Jersey, 1880–1920*. Chapel Hill: University of North Carolina Press.
- Roberts, Russell, and Rich Youmans. 1993. *Down the Jersey Shore*. New Brunswick, NJ: Rutgers University Press.
- Smith, Rogers M. 1992. "If Politics Matters: Implications for a 'New Institutionalism.'" *Studies in American Political Development* 6:1–36.
- Smith, Thomas F. X. 1982. *The Powerticians*. Secaucus, NJ: Stuart.
- Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge. 2015. "What Are We Weighting For?" *Journal of Human Resources* 50 (2): 301–16.
- Stokley, Alfred T. 2014. *Berkeley Township*. Charleston, SC: Arcadia.
- Tarrow, Sidney. 1998. *Power in Movement*. 2nd ed. Cambridge: Cambridge University Press.
- Vecoli, Rudolph J. 1965. *The People of New Jersey*. Princeton, NJ: Van Nostrand.
- Weisberg, Sanford. 1985. *Applied Linear Regression*. 2nd ed. New York: Wiley.
- Wilson, Harold F. 1953. *The Jersey Shore: A Social and Economic History of the Counties of Atlantic, Cape May, Monmouth and Ocean*, vol. 2. New York: Lewis Historical.
- Wortman, C. Byron. 1963. *The First Fifty: A Biographical History of Seaside Heights, New Jersey*. Seaside Heights, NJ: Fiftieth Anniversary Committee.