

The Insufficiency of “Democracy by Coincidence”: A Response to Peter K. Enns

Martin Gilens

In “Relative Policy Support and Coincidental Representation,” Peter Enns (2015) focuses on a crucially important question: Given the large disparities in political influence of more and less well-off Americans, “why don’t those in the economic middle elect politicians who might better represent their interests?” (1053).

This question has been debated at least since Werner Sombart asked why there is no socialism in the United States,¹ and the answer Enns proposes is not unlike Sombart’s. “We do not observe a political backlash from those in the economic middle,” Enns argues, “because policy typically corresponds with the median’s preferences” (1054). By this account, middle class Americans either don’t notice or don’t care about their lack of influence, because they are more or less satisfied with the policy outcomes that are determined by others.

Enns’ hypothesized explanation rests on two important points that I have repeatedly noted in my work on representational inequality: (1) the theoretical distinction between a group’s *influence* over government policy and the *congruence* of the group’s preferences with policy outcomes, and (2) the empirical observation that the policy preferences of different income groups are positively correlated and frequently aligned. Reflecting the combination of these two factors, Gilens and Page write “our evidence does *not* indicate that in U.S. policy making the average citizen always loses out. Since the preferences of ordinary citizens tend to be positively correlated with the preferences of economic elites, ordinary citizens often win the policies they want, even if they are more or less coincident beneficiaries rather than causes of the victory.”²

Enns and I agree that congruence without influence—what Gilens and Page called “democracy by coincidence”³—has important political implications. If policy outcomes were less congruent with the preferences of middle-class Americans than they currently are, we would indeed expect average voters to be less satisfied with their government and more cynical about the distorting power of moneyed interests than is already the case.

But there the agreement ends. In this paper, I spell out the problems with Enns’ critique. First, even if policy “typically corresponds with the median’s preferences,” as Enns writes, the atypical cases are, in fact, important and highly salient issues on which the power of the affluent and interest groups has pushed policy away from the preferences of the majority. There simply is not enough coincidence of policy outcomes and middle-class preferences to justify the conclusion that middle-income Americans are likely to be satisfied with the policies their government adopts. Second, Enns’ concept of “relative policy support” is vague and poorly developed. Third, Enns overlooks important characteristics of my data and makes erroneous calculations that undermine his conclusions. Finally, “democracy by coincidence” is not, in my view, a satisfactory substitute for true democratic responsiveness to the preferences of the governed.

Policy Congruence

The key empirical question raised in Enns’ critique is how congruent (or incongruent) policy outcomes are with middle-class preferences. Such incongruencies could arise for a variety of reasons. For example, a strong status quo bias lowers the probability of adoption even for policies that are popular at all income levels. But lack of responsiveness to the middle-class could also generate such policy incongruities if middle-class preferences diverge from the preferences of higher income Americans who are able to influence government policymaking.

The clearest indication of policy congruence (or its absence) arises from those proposed changes with strong majorities of middle-class Americans in favor or opposition. (In contrast, policy outcomes for which middle-income preferences are fairly evenly divided can tell us little

Martin Gilens is Professor of Politics at Princeton University (mgilens@princeton.edu). His research examines representation, public opinion, and mass media, especially in relation to inequality and public policy. Professor Gilens is the author of Affluence & Influence: Economic Inequality and Political Power in America (2012, Princeton University Press). He thanks Larry Bartels, Kosuke Imai, Benjamin Page, and Tali Mendelberg.

doi:10.1017/S1537592715002327

© American Political Science Association 2015

about congruence, since any set of outcomes on such policies would be about equally congruent with middle-class preferences.)

Looking at the same subset of 322 proposed policy changes that Enns analyzes,⁴ and defining “popular” policies as those that are favored by at least three-quarters of middle-income respondents and “unpopular” policies those that are opposed by at least three-quarters, we find congruence on only 34 percent of these proposed changes. (Reflecting the strong status quo bias evident in my data, policies that were unpopular with median income Americans failed 68 percent of the time, but policies that were popular with this group were adopted only 19 percent of the time.)

In contrast to the 34 percent “congruence rate” of the middle-class for these strongly favored or opposed policies, affluent Americans got the outcomes they preferred on 66 percent of the policies they strongly favored or opposed.⁵ (Again, a strong status quo bias is evident: 94 percent of the policies unpopular among high-income Americans failed, while 46 percent of popular policies were adopted.) Middle-class Americans, then, not only experience a low absolute level of policy congruence but are half as likely as affluent Americans to get the outcomes they prefer on the policies they strongly favor or oppose.

As these very different levels of policy congruence suggest, the policies that are popular (unpopular) among the well-off are not, for the most part, the policies that are popular (unpopular) among the middle-class. For example, of the 43 policies that are popular among middle-income respondents, only 8 are also popular among the affluent. Similarly, of the 19 policies that are unpopular among middle-income respondents, only 3 are also unpopular among the affluent.

What are these policies that are popular with the middle-class but not the affluent? The majority are redistributive policies including raising the minimum wage or indexing it to inflation, increasing income taxes on high earners or corporations, or cutting payroll taxes on lower income Americans. Other policies popular with middle-income but not high-income Americans involve greater government regulation of markets like restricting Japanese imports to help open the Japanese market to American goods, or preventing airlines from using bankruptcy laws to cut wages and break their unions. Policies unpopular with the middle-class but not the affluent tend to reflect the same preference divergences but in reverse. For example, middle-income Americans oppose tax cuts for upper-income individuals, spending cuts in Medicare, and roll-backs of federal retirement programs.

It would not be surprising if middle-class citizens noticed the failure of the federal government to raise the minimum wage, or the shifts in federal tax policy that primarily benefit the well off, or the adoption of trade

agreements that benefit American corporations but leave many American workers less well off. Having noticed these policy outcomes, it would not be surprising if these same citizens concluded both that people like them lack influence over what government does and that government policy has not been consistent with their interests and preferences, at least on a set of important and highly visible issues with strong middle-class majorities in favor or opposed.

Relative Policy Support

Enns begins his critique by focusing on congruence between middle-class preferences and policy outcomes, but then turns to a new approach to analyzing preferences that he calls “relative policy support.” Enns is unclear about just how relative policy support is defined, and he offers no formula by which it might be calculated. From his examples, it appears that relative policy support is intended to reflect the percentage point difference in support for two policies for a given group, and that this relative support is then to be incorporated in a comparison of policy preferences across groups. Enns’ figure 2 offers one example: if the percent favorable for policy 1 is 12 points higher than the percent favorable for policy 2 for group A, and also 12 points higher for group B, then groups A and B are said to have equal relative policy support regardless of the absolute levels of support on this pair of policies for each group.


How two groups might be compared when more than two policies are involved is unexplained, however, and some simple examples suggest that the difficulties are substantial. With three policies, there are three different policy pairs (policy 1 versus 2, 2 versus 3, and 1 versus 3), with 10 policies there are 45 pairs, and with the 322 policies Enns analyzes, 51,681 pairs. How is the degree of relative policy support to be aggregated across the multiple pairs, especially when the two groups will sometimes agree and sometimes disagree on which of the two policies in a pair is the more preferred?

If countervailing preferences are allowed to cancel out (for example, by averaging the sometimes positive and sometimes negative gaps in relative policy support), then very different patterns of preferences would produce the same overall “relative policy support” score. For example, compare the top two graphs in my figure 1. The top graph shows three policies on which group A is consistently 20 percentage points more favorable. Paralleling Enns’ figure 2, the relative policy support for each of the three pairs of policies in this top graph are all equal. (Groups A and B both prefer policy 2 over policy 1 by 5 percentage points, policy 3 over policy 1 by 20 percentage points, and policy 3 over policy 2 by 15 percentage points. The difference in relative support for each policy pair is zero.)

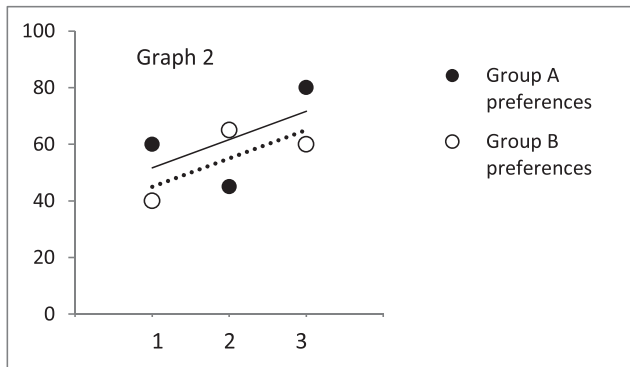
The middle graph in my figure 1 switches the support of groups A and B on policy 2. This scenario generates one

Figure 1


Three alternative hypothetical sets of three policies with varying levels of support from two groups; with the mean difference in relative support for the three policy pairs and the mean of the absolute difference in relative support for the three pairs



Policy pair	Group A Relative support	Group B Relative support	Difference in relative support
2 vs 1	+5	+5	0
3 vs 1	+20	+20	0
3 vs 2	+15	+15	0
mean difference:			0
mean absolute difference:			0



2 vs 1	-15	+25	+40
3 vs 1	+20	+20	0
3 vs 2	+35	-5	-40
mean difference:			0
mean absolute difference:			26.7



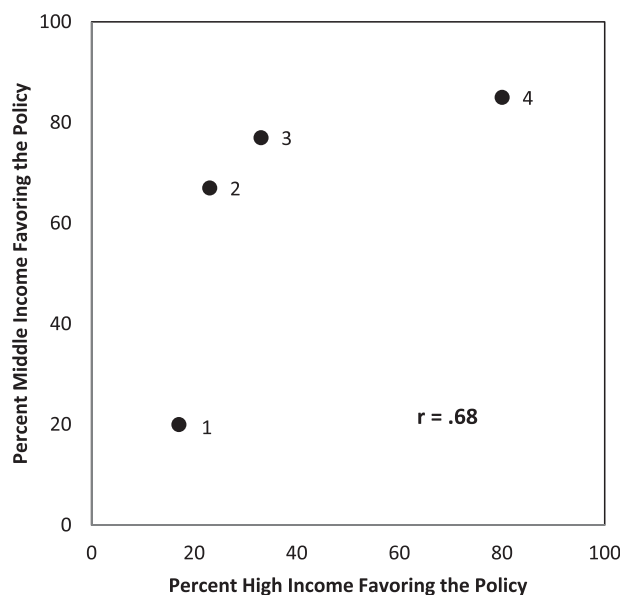
2 vs 1	+25	-15	-40
3 vs 1	+40	0	-40
3 vs 2	+15	+15	0
mean difference:			-26.7
mean absolute difference:			26.7

policy pair with equal relative support (both groups prefer policy 3 over policy 1 by 20 percentage points). But the other two policy pairs show opposite rankings for groups A and B: group A prefers policy 1 to policy 2 and policy 3 to policy 2 while group B prefers policy 2 to policy 1 and policy 2 to policy 3. Allowing these countervailing instances of relative support to cancel each other out would result in the identical (equal) overall relative policy support that the consistent pairs of policies in the top graph produced. It makes little sense to consider these two scenarios as similar in any comparison of policy support, given that one reflects consistent rankings on all three

policy pairs and the other reflects a consistent ranking on only one of the three policy pairs.

An equally problematic result would emerge if countervailing policy gaps were not allowed to cancel out. For example, if we take the mean of the *absolute* difference in “relative policy support” for each pair of policies, the middle and bottom graphs in my figure 1 would result in the same overall relative policy support score (+26.7). In each of these two graphs, there is one pair of policies with consistent relative support among groups A and B and two pairs with opposite and countervailing preferences. But *which* policies show agreement across groups

Figure 2
Hypothetical scenario in which strongly correlated preferences coexist with strongly divergent levels of support for policies 2 and 3



and which show disagreement matters. If policies with greater support are more likely to be adopted, then group A in figure 3 will enjoy much more congruence and a much stronger association between preferences and outcomes, than group B.

Neither of these alternative methods of calculating relative policy support is able to capture the meaningful differences in the preference patterns across groups in these three hypothetical scenarios. And this problem is anything but hypothetical. Across the 51,681 possible pairs of the 322 proposed policy changes that Enns analyzes, middle- and high-income Americans ranked the paired policies in the opposite direction 24 percent of the time. That is, in about one case out of four, the policy that was more preferred by the middle class was less preferred by the affluent and vice versa.

These examples are not intended to explicate Enns' half-articulated conception of relative policy support, but only to underline one of the difficulties that would have to be resolved before it could become a conceptually and empirically useful idea.

Coincidental Representation

When Enns moves from his conceptual discussion to his empirical analysis of my data, he shifts his focus from relative policy support to the correlation of preferences across income groups (a related but distinct concept). As I've noted in my previous publications on representational inequality, preferences across income groups are positively

correlated; *on average* policies that are more popular with the affluent are also more popular with the middle class and the poor. This is true for the full set of 1,779 policies in my 1981–2002 dataset and for the 322 policies on which preferences of the fiftieth and ninetieth income percentiles diverge by at least 10 percentage points (i.e., the set of policies Enns analyzes).

But even a strong correlation between two groups' preferences need not imply similar levels of congruence between preferences and outcomes. In my figure 2, support for the four policies is correlated at .68 (identical to the actual correlation between middle- and high-income preferences among the 322 proposed policy changes Enns examines). While both groups favor policy 4 and oppose policy 1, policies 2 and 3 are favored by the middle-class but opposed by the well-off. Despite the seemingly strong correlation, the congruence between preferences and outcomes would depend strongly on which group policymakers responded to.

After noting the positive correlation between middle-class and affluent preferences in my data, Enns turns his attention back to policy congruence but with a weaker standard. Rather than claiming that policy outcomes “typically correspond with the median's preferences,” he now claims that “policy ends up about where those in the middle would expect if they received the same representation as affluent individuals” (1057).

I showed earlier that this is not true for policies that are favored or opposed by at least 75 percent of middle-income Americans: the middle-class and the affluent differ on which policies they strongly favor and oppose, and the affluent are much more likely to see their preferences on such policies reflected in government decision making.

Enns' analyses on this point are displayed in his figure 6. These results do not convey the consequences of representational inequality for three reasons. First, Enns shows the predicted outcomes for both the fiftieth and ninetieth income percentiles (on the Y axis) in relation to the observed policy preferences of the ninetieth income percentile (on the X axis).⁶ What is not evident in the figure, but important in understanding the consequence of representational inequality, is that the set of policies that would fall at the high and low ends of the preference distribution would differ if we were considering a scenario in which the middle-class rather than the affluent shape policy outcomes.

The second reason that Enns' figure 6 gives a misleading impression is that the functional form imposed in generating predicted probabilities is a good fit for the full set of proposed policy changes⁷ but among the subset of 322 policies on which preferences of the fiftieth and ninetieth percentiles diverge, the model over-predicts the probability of adoption for policies that are strongly opposed by the affluent or strongly favored by the middle-class.⁸ Consequently, the hypothetical world in which the middle-class has the degree of influence over policymaking that the

affluent now enjoy would differ more from actual conditions than the model underlying figure 6 suggests (at least for those policies with especially high or low levels of support).

Third, in figure 6 (and the analogous figures in his appendix), Enns errs in calculating the statistical significance of the difference between the two sets of predicted probabilities. Describing the results underlying figure 6, Enns claims that “approximately 66 percent of the predicted values based on middle-income preferences fall within the 95 percent confidence intervals of the predicted values for the affluent” and that “none of the predicted values for the middle-income are statistically different from the affluent” (1058). However, assessing the statistical significance of the difference between the two predicted probabilities of adoption for each of the 322 proposed policy changes is more complex than Enns’ analysis suggests.

First, these comparisons result from substituting alternative preference scores into the same model based on the link between preferences and policy outcomes for the ninetieth percentile. What Enns fails to note in his assessment of statistical significance is that the model-based errors of prediction estimated from affluent and middle-class preferences are perfectly correlated (because they are based on the same regression model). The standard errors Enns shows in his figure 6 result from uncertainty about the true model (based on the desire to generalize the results from the 322 observed policies to a larger hypothetical population of similar policies). But if the model-based predictions for the ninetieth percentile are too high or low at a given level of favorability, that model-based error will apply to the predictions of the fiftieth percentile as well. Indeed, if the slope of the estimated association is significantly different from zero, and the preferences of the fiftieth and ninetieth percentiles diverge for a particular policy, then the two predicted probabilities of adoption must differ significantly from each other (and, in fact, the slope of the line in Enns’ figure 6 is significant at $p < .01$ with a two-tailed test).

Enns also ignores the fact that the preferences of the fiftieth and ninetieth percentiles in his calculations are themselves uncertain estimates. However, this additional source of statistical uncertainty turns out to be much too modest to rescue his claim that “none of the predicted values for the middle-income are statistically different from the affluent.” The estimated preferences at these two income levels are imputed by regressing support or opposition to the policy change on survey respondents’ income and income-squared, separately for each of the 322 policies.⁹ Calculating the significance of the difference between the imputed preferences for each policy is complex, but a maximum limit can be estimated by treating the two groups’ imputed preferences as statistically independent and using the standard errors of each of the

two imputed preferences to calculate the standard error of the difference between them.¹⁰ Doing so shows that the imputed preferences of the fiftieth and ninetieth income percentiles differ at $p < .05$ for 307 of the 322 policies and at $p < .01$ for 292 of them.

In short, simply observing the confidence intervals of the predicted probabilities for the ninetieth income percentile as shown in figure 6 (or the confidence intervals for both the fiftieth and the ninetieth percentiles), cannot reveal whether the differences in probabilities between the fiftieth and ninetieth percentiles are significant. A proper accounting of the two sources of sampling error in these estimates, however, reveals low levels of statistical uncertainty regarding both the positive slope of policy outcomes on policy preferences (for the ninetieth percentile) and, in the vast majority of cases, the difference between the imputed preferences for the fiftieth and ninetieth percentiles.

Strong Partisans’ Preferences

As Enns notes, income-based political inequality is not the only potentially important dimension of inequality we might care about. Enns illustrates the potential for large average differences in policy preferences by focusing on strong Democratic and strong Republican identifiers. These groups’ preferences are indeed strongly divergent, although that is hardly surprising given the polarized structure of American politics and Enns’ choice of groups at the two ends of the political spectrum.

Examining outcomes, Enns finds that in a five year period in which “national politics strongly favored Republicans” (1059), policy outcomes reflected the preferences of strong Republicans. Because strong Democrats’ and strong Republicans’ preferences are only weakly correlated, strong Democrats did not receive much “coincidental representation” of the sort middle-class Americans get. Again, as Enns acknowledges, this is hardly surprising.

But unequal responsiveness based on partisan preferences can coexist with unequal responsiveness based on income. Bartels, for example, shows dramatic differences in the voting behavior of Democratic and Republican Senators, but also shows that neither Democratic nor Republican Senators were at all responsive to the preferences of their least well-off constituents, and that Republican Senators were about twice as responsive to their high-income as to their middle-income constituents.¹¹

From a normative standpoint, partisan representation does not constitute the challenge to American democracy posed by income-based representational inequality. Each political party is expected to “represent” its partisan constituents and pursue the policies they prefer; indeed, that is a (if not the) central function of parties in a democracy. In a period in which one party is strongly favored by its dominance of national political institutions, policy outcomes in a well-functioning democracy should

reflect the policy commitments of that party, at least to a substantial degree. In contrast, my analyses of many decades of policy change under a wide range of political circumstances never revealed a period in which middle-income Americans had as much influence over government policy as the affluent, and rarely had any apparent influence at all.

The Merely Affluent and the Truly Rich

Another consideration in assessing the implications of representational inequality for policy congruence stems from an important limitation of my data. As I’ve argued elsewhere, analyses of the ninetieth income percentile likely understate the true power of the rich to shape policy, and consequently the true impact of representational inequality.¹² This is because influence among high income Americans is not equally shared: those at the 99th percentile surely exert more influence over government policymaking than those at the ninetieth, and those at the 99.9th even more.

The available data on the preferences of the truly rich suggest that their views tend to differ from the median income American’s along the same lines, but to a larger degree, than those of the ninetieth income percentile.¹³ If political power is disproportionately concentrated among the top one percent, or the top one-tenth of one percent, then the apparent degree of coincidental representation evident in my dataset (and Enns’ analyses) is exaggerated. First, because the preferences of the relevant income groups differ more than my data for the ninetieth income percentile suggest, and second, because the influence of the rich is underestimated when using my ninetieth income percentile proxy (and if the influence of the rich is underestimated, then the alternative scenario in which this degree of influence is “assigned” to the middle-class understates the consequence of class differences in policy preferences).

“Democracy” by Coincidence?

Enns suggests that the lack of political backlash among middle-income Americans arises from the correspondence between their preferences and actual policy outcomes. It is certainly true that under the right circumstances, a powerless group can enjoy the same policy outcomes that actual influence would produce. In an absolute theocracy, for example, any group of citizens that shares the preferences of the theocratic rulers will get the policies they desire.¹⁴

But democracy by coincidence is a debased and conditional form of democracy (if it is a form of democracy at all). If the majority of citizens must depend on their agreement with a powerful minority in order to obtain the policies they prefer, that majority will always be in a politically tentative and precarious position. As I argued earlier, middle-income Americans’ preferences frequently differ from the policies their government has

pursued over the past decades. But even if they did not, the majority of Americans might well be dissatisfied with a political system that treats them as spectators rather than citizens.

Conclusion

Peter Enns is quite right that “democracy by coincidence” is an important feature of contemporary American politics. And he is also right that the greater the congruence between policy outcomes and middle-class preferences, the less politically dissatisfied middle-income Americans are likely to be. But his conclusion that this coincidental representation is strong enough to generate satisfaction with current policy among middle-income Americans is, I believe, mistaken.

On those policies where strong congruence or incongruence is possible—that is, policies with either strong support or opposition—middle-class preferences tend not to prevail. Of course, affluent Americans do not always get the policies they prefer either. But the affluent are twice as likely to see the policies they strongly favor adopted, while the policies they strongly oppose are only one-fifth as likely to be adopted as those that are strongly opposed by the middle class.¹⁵ On important aspects of tax policy, trade policy, and government regulation, both political parties have embraced an agenda over the past few decades that coincides far more with the economically regressive, free trade, and deregulatory orientations of the affluent than with the preferences of the middle-class.

Unequal influence does matter, first, because coincidental representation is a pale, counterfeit, simulacrum of democracy. And second, because in its current form, policy outcomes as shaped by the well-off are incongruent with the preferences of the middle-class on too many salient and important policies.

Notes

- 1 Sombart 1906.
- 2 Gilens and Page 2014, 572.
- 3 Ibid., 573.
- 4 These include the subset of proposed policy changes on which preferences of the fiftieth and ninetieth income percentiles differ by at least 10 percentage points.
- 5 Using the same criteria of under 25 percent or over 75 percent favorability. The differing congruence rates for middle- and high-income Americans cannot be accounted for by different levels of favorability within the groups of strongly favored and opposed policies; the mean favorability among strongly opposed policies is slightly lower for the fiftieth percentile than the ninetieth percentile (18 versus 20 percent), and slightly higher for strongly favored policies (82 versus 80 percent).
- 6 In Enns’ appendix, a parallel chart is provided to that in figure 6, but with both sets of predicted outcomes

- arrayed along an X axis that represents support among the middle class. Neither of these charts reveals the extent to which popular (or unpopular) policies differ between the two groups.
- 7 Gilens 2012, 75.
 - 8 Specifically, the model predicts that 18 percent of policies unpopular with the affluent would be adopted compared with the observed 6 percent (counting policies with less than 25 percent support as unpopular). Among policies popular with the middle-class (i.e., with greater than 75 percent support), the model predicts 31 percent would be adopted compared with an observed 19 percent.
 - 9 Gilens 2012, 61.
 - 10 The variance of the difference between two random variables is equal to the sum of their variances minus twice their covariance. Consequently, the true standard error of the difference between the (positively covarying) imputed preferences for the fiftieth and ninetieth income percentiles is smaller than the estimate that would be produced by treating these two imputed values as statistically independent (and thereby ignoring their positive covariance).
 - 11 Bartels 2008.
 - 12 Gilens 2012, 242; Gilens and Page 2014.
 - 13 Page, Bartels, and Seawright 2013; Gilens and Page 2014.
 - 14 Gilens 2012, 66.
 - 15 Including policies that were strongly favored (over 75 percent) or strongly opposed (under 25 percent) by both groups, 46 percent of those strongly favored by the affluent were adopted compared with 19 percent for the middle-class; six percent of those strongly opposed by the affluent were adopted compared to 32 percent for the middle-class.

References

- Bartels, Larry M. 2008. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton, NJ: Princeton University Press.
- Gilens, Martin. 2012. *Affluence and Influence: Economic Inequality and Political Power in America*. Princeton, NJ, and New York: Princeton University Press and the Russell Sage Foundation.
- Gilens, Martin and Benjamin I. Page. 2014. "Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens." *Perspectives on Politics* 12(3): 564–581.
- Page, Benjamin I., Larry M. Bartels, and Jason Seawright. 2013. "Democracy and the Policy Preferences of Wealthy Americans." *Perspectives on Politics* 11(1): 51–73.
- Sombart, Werner. 1906. *Warum gibt es in den Vereinigten Staaten keinen Sozialismus?* Tübingen: J. C. B. Mohr (Paul Siebeck). Republished 1976 as *Why Is There No Socialism in the United States?*, London, Macmillan.