

Testing Inferences about American Politics: A Review of the “Oligarchy” Result

Research and Politics
 October-December 2015: 1–7
 © The Author(s) 2015
 DOI: 10.1177/2053168015608896
 rap.sagepub.com


Omar S Bashir

Abstract

In a well-publicized study, Gilens and Page argue that economic elites and business interest groups exert strong influence on US government policy while average citizens have virtually no influence at all. Their conclusions are drawn from a model which is said to reveal the causal impact of each group’s preferences. It is shown here that the test on which the original study is based is prone to underestimating the impact of citizens at the 50th income percentile by a wide margin. In addition, descriptive analysis of the authors’ dataset reveals that average Americans have received their preferred policy outcome roughly as often as elites have when the two groups have disagreed with each other. Evidence that average citizens are effectively ignored by the policy process may not be as strong as is suggested by the authors.

Keywords

American politics, representation, inequality, quantitative methods, replication

Introduction

In their influential article entitled “Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens,” Martin Gilens and Benjamin I. Page use statistical analysis to adjudicate between four ideal-type theories of American politics (Gilens and Page, 2014). Their main findings are that “economic elites and organized groups representing business interests have substantial independent impacts on U.S. government policy, while average citizens and mass-based interest groups have little or no independent influence.” These findings provide support for theories called economic-elite domination and biased pluralism.

According to several journalistic accounts but not Gilens and Page themselves, the findings show that the American system of government is best understood as “oligarchy.”¹ Extensive press coverage of the article has successfully drawn attention to one of the most important questions in the study of contemporary American politics: to what extent do the wealthy dominate average citizens in the formulation of government policy?² In pursuit of an answer, Gilens and his team of researchers gathered data over a long period, tracking 1923 instances between 1981 and 2002 in which national surveys asked favor/oppose questions about proposed policy changes. This is

a commendable attempt to bring evidence to the study of a timely and politically-loaded topic. Yet, as Gilens notes in his 2012 book based on a similar empirical foundation, even the most meticulously assembled dataset may not lend itself to straightforward inference.

Gilens and Page nevertheless draw several strong inferences from their analysis. The authors do not argue that policy outcomes correspond disproportionately to the preferences of the wealthy; in fact, their dataset reveals that the wealthy and the average have highly correlated preferences. Rather, the main inferences in the paper are about causality: do we only have “democracy by coincidence” in the United States? Gilens and Page claim that they can “decisively reject” majoritarian views of American democracy because they have found a way to determine which groups have “independent influence” in policymaking.³ According to the authors, “[o]ur main point concerns causal inference: if interpreted in terms of actual causal impact,

Princeton University, USA

Corresponding author:

Omar S Bashir, Department of Politics, Princeton University, 130
 Corwin Hall, Princeton, NJ 08544, USA.
 Email: obashir@princeton.edu



Creative Commons Non Commercial CC-BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 3.0 License (<http://www.creativecommons.org/licenses/by-nc/3.0/>) which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the SAGE and Open Access pages (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

Table 1. Reprinted from Table 3 in Gilens and Page. From the original caption: the dependent variable is the policy outcome, coded 1 if the proposed policy change took place within 4 years of the survey date and 0 if it did not. Predictors are the logits of the imputed percentage of respondents at the 50th (“average citizens”) or 90th (“economic elites”) income percentile that favor the proposed policy change... All analyses reflect estimated measurement error in the predictors.

	Model 1	Model 2	Model 3	Model 4
Preferences of average citizens	0.64***	—	—	0.03
Preferences of economic elites	—	0.81***	—	0.76***
Alignment of interest groups	—	—	0.59***	0.56***

*** $p < 0.001$, $N = 1779$.

the prior findings [supporting majoritarian theories] appear to be largely or wholly spurious.”⁴ They point to a “nearly total failure” of majoritarian frames and assert that “the preferences of economic elites... have far more independent impact upon policy change than the preferences of average citizens do.”⁵ They conclude that “America’s claims to being a democratic society are seriously threatened.”⁶

After summarizing the steps taken by Gilens and Page, I examine the statistical basis for their central claim that average Americans have virtually no influence on policy outcomes. I show that the result on which the original study is based is too likely to have been produced by chance because the income-based independent variables are highly correlated. I then evaluate three of the study’s descriptive claims about American democracy before concluding.

Summary of original approach

This section outlines the steps taken to reproduce Table 3 in the original paper. Reprinted here as Table 1, it features the main result of the study. Several predicted probability plots and odds comparisons in the original article are based on the coefficients in Model 4 of the table. Gilens and Page show in the first three columns that each of their independent variables, one at a time, seem to exert a positive and significant effect on policy outcomes.⁷ These independent variables include the preferences of average citizens (proxied by the estimated preferences of respondents at the 50th income percentile), the preferences of economic elites (90th income percentile), and the preferences of interest groups.

“But the picture changes markedly,” the authors state, “when all three independent variables are included in the multivariate Model 4... The estimated impact of average citizens’ preferences drops precipitously, to a non-significant, near-zero level. Clearly the median citizen or ‘median voter’ at the heart of theories of Majoritarian Electoral Democracy does not do well when put up against economic elites and organized interest groups.”⁸ Note that the authors’ basis for causal inference is the inclusion of multiple variables in the same model.

Model 4, the key result in the paper, does not reflect the output of a typical logistic or linear regression test. The authors pursue a rather non-standard approach because they

identify correlated survey error between the 50th and 90th income percentile preference variables. As they explain in their Appendix 2, a typical multiple regression⁹ including both those variables produces implausible coefficient estimates. These implausible estimates are attributed to the correlated survey error just mentioned.¹⁰

Given their diagnosis, Gilens and Page perform a multi-step correction procedure. First, they quantify correlated survey error by exploiting groups of two, three, or more of the 1779 survey questions which they code as addressing the same basic concept in the same calendar year. The authors claim that, based on their identification of 116 sets of similar survey questions, measurement error is responsible for 17% of the observed covariance between the measured 90th and 50th percentile variables.¹¹

Next, using that estimate of correlated measurement error, the authors “estimated structural equation models in AMOS that purged of error the structural coefficients representing the associations of the predictors with [their] outcome measure.” The structural equations are not specified in the paper. AMOS features a graphical interface in which users draw model diagrams. Unfortunately, the input diagram for the central statistical test is not included in the paper’s replication repository,¹² but one of the authors kindly shared it with me. The input diagram consists of two linked components as in Arbuckle (2012).¹³ One is a measurement submodel which accounts for the aforementioned correlated error between the income-based variables and also for measurement error in the interest-group variable alone (the authors estimate a reliability of 0.87 for that measure). Second, the core submodel linking the unobserved but “corrected” versions of independent variables to the dichotomous dependent variable is a linear regression model for which AMOS provides coefficients. These coefficients are those reported in Model 4.

For presentational purposes, the authors use the coefficients to compute predicted probabilities of policy change. The plots in the study reveal a flat line of virtually zero policy responsiveness to the preferences of average Americans. Meanwhile, elites’ preferences seem to swing the predicted probabilities of policy change dramatically. I was unsure about how to repeat the process of computing those probabilities since Model 4 is a modified linear

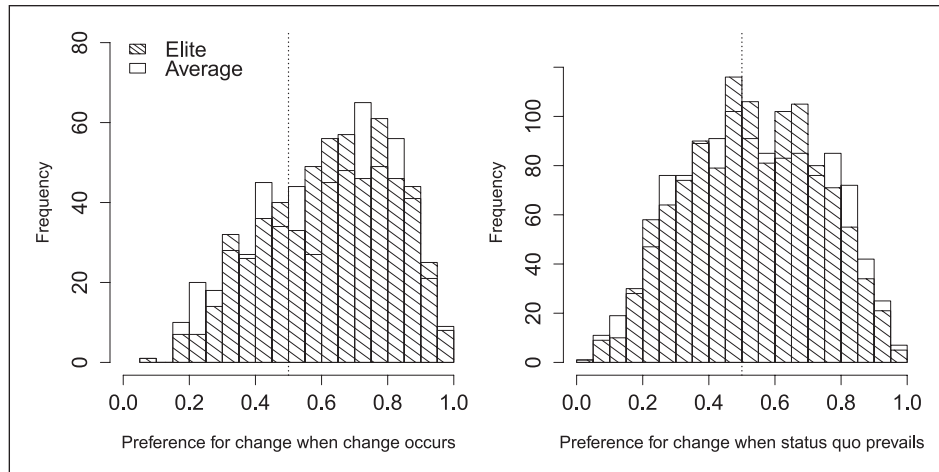


Figure 1. Histograms which illustrate the extent of overlap in the preference distributions of 90th and 50th income percentile Americans, conditional on policy outcome. Left: only those 588 cases in which change ensues. Right: the 1191 cases in which no full policy change is detected within 4 years.

regression, not a logistic regression. But because the predicted probability plots are all based on the coefficients reported in Model 4, any problems with the divergent coefficients representing the median- and high-income groups would implicate the plots as well. The next section thus asks whether there is an alternative explanation for the difference in reported coefficients.

Deceptively divergent income coefficients

Recall from Table 1, Model 4 above the striking difference reported between income-based coefficients. The coefficient for 90th income percentile Americans is highly significant ($p < 0.001$). At a value of 0.76, it is over 25 times larger than the coefficient of 0.03 corresponding to the median-income group. Because the 50th percentile coefficient takes on a “near zero” value, the authors claim that the policy process is non-responsive to average Americans.

But the original study relies on linear regression of a dichotomous dependent variable on two highly correlated independent variables. Standard practice when dealing with dichotomous outcomes is to employ logistic regression to avoid violation of the constant error variance assumption of linear regression. In addition, high correlation between independent variables violates an assumption of both linear and logistic regression. The correlation coefficient between the income-based variables is $r = 0.78$, even after the authors’ procedure to address correlated error reduces the coefficient from its observed value of $r = 0.94$. Another reason to investigate further is that the preference distributions of the two income groups are difficult to distinguish from each other when conditioned on policy outcome, as shown in Figure 1. The high correlation between variables already

suggests that the unconditioned distributions are similar to each other. Yet if one group is far more influential on policy, one might expect distributions to diverge when conditioning in this way. Group means are within 0.02 of each other in both plots.

I employ a simulation to investigate whether the authors’ linear regression of a dichotomous dependent variable on highly correlated independent variables can generate extreme but incorrect results. Online Appendix A lists the specific steps in the simulation, which is implemented in R with replication code accompanying this review.¹⁴

While the steps are detailed elsewhere, here I provide an overview of the procedure. For each simulation iteration, I randomly generate three independent variables, ensuring that they each have the properties and mutual relationships that the Gilens and Page variables have after the authors’ error corrections are applied. This means, for instance, ensuring that two of the randomly generated variables are highly correlated ($r = 0.78$). Next, I construct an outcome variable for the iteration by first choosing a “true” coefficient for each independent variable and then using a linear model to compute the left-hand side in the familiar regression equation setup based on those coefficients. The outcome variable is dichotomized to match the form of its analogue in the study. Finally, I perform linear regression of the outcome variable back on the randomly generated independent variables, which should yield estimated coefficients close to the three true coefficients that defined the data-generating process. If linear regression fails to produce accurate estimates of the true coefficients, this signals a problem with the numerical conditions in the study. Specifically, if I choose a true coefficient for the analogue to the median-income independent variable that is reasonably large, but the approach tends to produce a much smaller and thus erroneous estimate along with high levels

of apparent statistical significance, then the Gilens and Page result may not be reliable.

Another way to describe the simulation is that it computes regression coefficient *estimates* under a fixed set of true (chosen) coefficients. It allows one to ask, “If the true coefficients for the 50th and 90th percentile variables were in fact β_1^t and β_2^t , what kinds of estimates of those coefficients, $\hat{\beta}_1$ and $\hat{\beta}_2$, would the study’s approach tend to produce?” The chosen coefficients β_1^t and β_2^t are true in the sense that they are used in each simulation iteration to construct an outcome variable from randomly-generated versions of the independent variables, matching as closely as possible the conditions in the original study. One would expect the estimated coefficients produced by subsequent regression to be close to the true coefficients chosen to seed each iteration.

Note that I am not simply questioning the statistical significance of the difference between the 50th and 90th percentile coefficients, which would require only a test of the null hypothesis of coefficient equality. In other words, Gilens and Page do not merely argue that the income-based coefficients are statistically different from each other; indeed, prior research already holds that the wealthy have a moderately greater impact on policy outcomes. The authors go further by stressing a drastic substantive difference in coefficient magnitude, with one coefficient being virtually zero.

Results

What if the 50th percentile coefficient were in reality not minuscule? Would the study’s approach report that it was near zero nonetheless? I run multiple simulations to answer the question. Across simulations, I vary the true coefficient β_1^t corresponding to the median-income variable at values much larger than its reported value of 0.03. This allows one to examine the rate of erroneously small estimates as a function of the values that β_1 might actually take. I hold the true coefficient β_2^t corresponding to elites at its reported value of 0.76 in every simulation.

The results are displayed in Figure 2. If β_1^t is about 0.4, larger than half of the high-income coefficient, the statistical approach in the study mistakenly estimates it to be essentially zero in more than 20 percent of trials.¹⁵ We also see the study’s extreme divergence between $\hat{\beta}_1$ and $\hat{\beta}_2$ at a rate greater than 10 percent when the chosen coefficient is set to that value.

Not shown in the figure is that when the enforced correlation between independent variables is reduced from 0.78 (as in the study) to lower values, the simulation does not produce extreme results by either criterion, $\hat{\beta}_1$ near zero or wide divergence between $\hat{\beta}_1$ and $\hat{\beta}_2$, for any value of β_1^t tested. That reliable results are produced when this change is made confirms that the simulation is not stacked against the original approach.

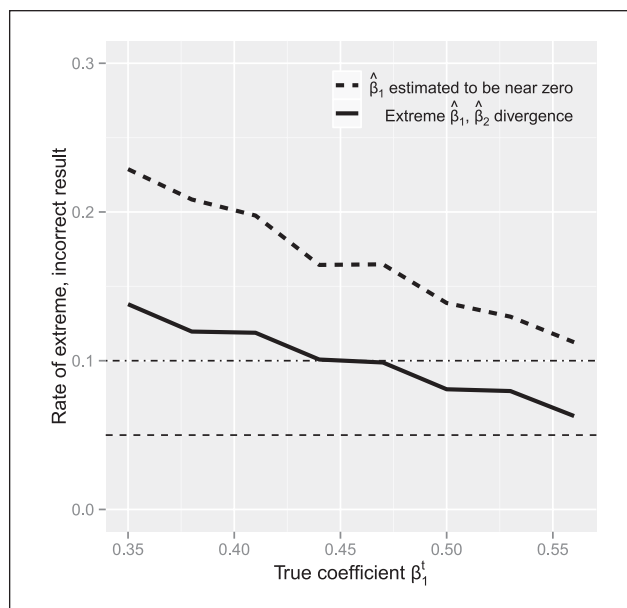


Figure 2. Simulation results. Whether one adopts a 90% or 95% confidence standard (horizontal lines), the authors’ method too often produces extreme but erroneous results. The method frequently yields a value of $\hat{\beta}_1$ near zero even when β_1^t , the true coefficient value (x-axis), is much higher. The two ways to interpret the study’s main result are listed in the legend. Near or “essentially zero” is defined by the authors to be $\beta_1 \leq 0.05$. “Extreme $\hat{\beta}_1, \hat{\beta}_2$ divergence” entails both $\hat{\beta}_1 \leq 0.03$ and $\hat{\beta}_2 \geq 0.76$, with the latter significant at the 99.9% level.

When we return to the study’s correlation level (Figure 2) and instead set β_1^t to 0.56, equal to the coefficient reported for interest groups, the estimated coefficient $\hat{\beta}_1$ is still found to be essentially zero in more than ten percent of simulation trials even though it is known to be much larger. The extreme divergence found in the study is also still produced at a rate that exceeds the common 95% significance standard. Even if the median-income and interest group coefficients were in fact equal to each other, the authors’ approach would too often produce the numbers they report simply by chance.

The main point is not a single significance test, but rather that Figure 2 illustrates how multivariate regression under these conditions is indeed prone to overstatement of the importance of income. Furthermore, it makes sense that $\beta_1^t = 0.03$ (not shown) is more likely than any of the larger chosen coefficients tested in the figure to yield the wide coefficient divergence (0.03, 0.76) reported in the study. But it yields that divergence and corresponding statistical significance at a rate of only 0.3. Thus, a 50th percentile coefficient of $\beta_1 = 0.41$ is roughly 40% as likely as the reported value of $\beta_1 = 0.03$ to have produced the main result on which the study is based.

Figure 3 illustrates another problem with the method’s performance. The coverage ratio for β_1^t , defined as the proportion of trials in which the 95% confidence interval

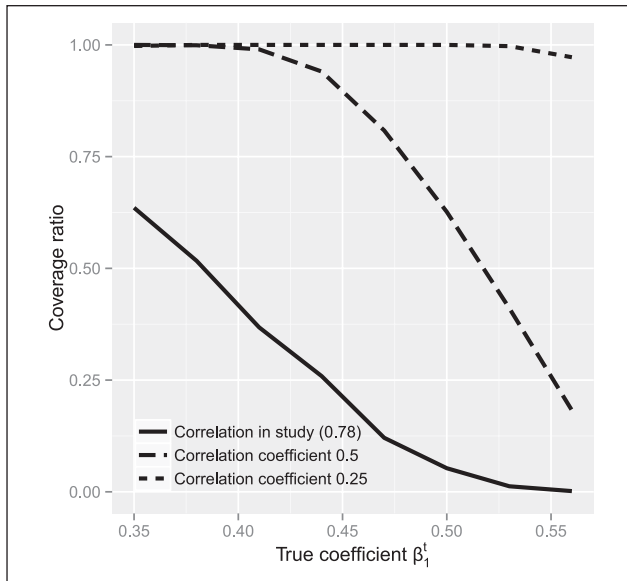


Figure 3. When independent variables are as highly correlated as they are in the study (solid curve), the 95% confidence interval frequently misses the true coefficient value (low coverage ratio). Performance of the study's approach improves at lower levels of correlation.

around the point estimate $\hat{\beta}_1$ contains the true value of the coefficient β_1^t , is very low for those simulations in which the two independent variables of interest are as highly correlated as those in the study. When the enforced correlation coefficient between the first two independent variables is reduced from 0.78 to lower positive values, the performance of multivariate analysis improves. Thus, even though standard errors are reported to be small, the high correlation between the preferences of median-income and high-income Americans seems to interfere with reliable estimation.

Online Appendix B gauges whether a simpler kind of test can support the assertion that average Americans have no impact on policy. Conditional on interest group opposition to change, the policy process seems more responsive to median-income citizens than to economic elites. On the other hand, the preferences of the wealthy seem to have much more impact when interest groups support change, though the number of cases involved in the analysis is small.¹⁶ Future research might investigate this divergence further. The possibility that differences in income-based responsiveness are conditional is another potential caveat to the study's findings.¹⁷

Distinguishing and evaluating other claims

The previous section shows that there is not yet enough evidence for the claim that average citizens have very little impact on policy outcomes. Yet there are three descriptive

claims in the original paper which should be examined as well. These assertions, some of which have been emphasized in popular discourse, are more distinct from the central claim and from each other than they at first appear to be.

First, the authors state that “even when fairly large majorities of Americans favor policy change, they generally do not get it.”¹⁸ Gilens and Page reference in their conclusion their descriptive finding that, even if 80% of the public favors change, that change occurs less than half of the time.¹⁹ Readers of the concluding section may not realize that “public” includes elites. In the original dataset, change is enacted 47% of the time that median-income Americans favor it at a rate of 80% or more. Yet change is enacted 52% of the time that elites favor it at that rate. The difference between groups is smaller when one examines not only strong preferences for change but strong preferences for either policy outcome.²⁰ The authors mention but do not emphasize that elites, too, seem to be affected by a status-quo bias. It is not clear how this finding is consistent with a story of elite domination, especially because average citizens tend to support the status quo more often when the groups disagree.

Second, Gilens and Page claim that “reality is best captured by” theories in which both economic elites and organized interest groups “play a substantial part in affecting public policy...”²¹ Shortly after claiming that their model captures reality, the authors caution that the R^2 value for Model 4 is 0.074. Roughly speaking, that means that their model, which accounts for the very groups that they say play a substantial, even dominant²² role in determining policy, explains less than 10 percent of the observed variation in policy outcomes. The drastically different coefficients (0.03 and 0.76) reported for the two income groups can be exchanged with each other and the resulting model still successfully predicts almost the same number of policy changes in the sample.²³

The authors acknowledge the low R^2 value and list potential reasons.²⁴ The low value is not necessarily a problem for hypothesis testing. Still, it provides useful information. For instance, when *neither* the rich nor the average favor change, change still happens at a rate of 23% in the dataset. The policy process seems only weakly responsive to the preferences of the wealthy compared to variables missing from the model.

Finally, Gilens and Page claim that ordinary citizens get the policies they favor “only because those policies happen also to be preferred by the economically-elite citizens who wield the actual influence.”²⁵ The authors' main focus is on causality, but they also make the *descriptive* claim that when average citizens disagree with elites or organized interests, “they generally lose.”²⁶ To commentators, this interpretation seemed to capture the essence of the project. One prominent voice summarized the entire study in the following way: “when elite preferences and popular preferences are different, the elite almost always wins.”²⁷

Yet, this is contradicted by the authors' dataset. There are 185 cases in the data in which the average preferences of the two income groups are on opposite sides of an issue.²⁸ Median-income Americans receive their desired outcome 47% of the time that the policy process must pick a winner between the average and the elite since the two groups disagree. The results are similar when the analysis is restricted to only those cases of disagreement which also exhibit a large preference gap between groups.²⁹ Nor do the results change if interest groups are incorporated as follows. The rich get their favored outcome despite the combined opposition of the other two groups at a rate of 32%; meanwhile, average Americans' favored outcome occurs 30% of the time that they face combined opposition from interest groups and the wealthy. It is true that median-income citizens are more likely to prefer the status quo when they and the wealthy disagree, but this suggests that any status-quo bias embedded in the policy process favors average Americans.

Conclusion

Even if I have not erred in this review, it would be wrong for readers to conclude that the wealthiest Americans and business interests do not enjoy advantages in influencing the policy process. The Gilens and Page (2014) article is only one part of a growing body of scholarship on this topic,³⁰ and further work may uncover evidence that these advantages are in fact overwhelming. In addition, even if inequality were somehow shown to have no bearing on who influences policy, it would still be morally wrong to ignore it.

Yet, what this review aims to highlight is that the original study exhibits weaknesses in its main causal claim and in three of its descriptive claims. The statistical approach employed in the study's central test seems too unreliable to gauge how much influence median-income citizens enjoy relative to elites and interest groups. The combination of a linear model, dichotomous dependent variable, and high correlation between independent variables yields misleading estimates. The coefficient representing the influence of median-income citizens could be as large as the coefficient for interest group influence. The more important issue is that the study's approach has poor resolution on the median-income coefficient. The approach also often produces confidence intervals which do not contain true coefficient values.

In short, the analysis is prone to underestimating drastically the causal impact of median-income preferences, assuming that regression coefficients even capture causality in this context: the authors' claim to causal inference is based only on the fact that they perform multiple regression. The authors have not yet shown that prior findings more amenable to majoritarian theories are "largely or wholly spurious," nor do their results seem to enable adjudication between competing conceptions of American democracy.

I also evaluated a set of secondary claims in the study. The notion that the American system is mere "democracy by coincidence" must contend with the finding that average

Americans have received their desired outcome roughly as often as the richest have when the two groups have been on opposite sides of an issue. Any status-quo bias in the policy process affects both income groups to a similar extent, and it may favor average citizens, who prefer the status quo more often in the data. In addition, the authors' model explains little of the variation in policy outcomes, so economic elites and interest groups cannot be said to "dominate" policymaking on the basis of this research even if they do have a greater impact than average citizens.

Although the authors' potentially fruitful distinction between different types of interest groups was not the focus of this review, the original study's result regarding the advantage of business groups over mass-based groups is unclear. Gilens and Page emphasize a much larger regression coefficient for business groups than for mass-based groups in their Table 4. Yet they then report that, after they adjust for the number of actors of each kind, the two interest group types have roughly equal influence. The predominance of business interest groups in the study thus rests on the fact that there are more of them included in the analysis, which is partly a result of the authors' choice to add business groups they felt were missing from *Fortune* magazine's "Power 25" lists.³¹ It is not clear, then, that mass-based groups (labor organizations but also, by the authors' definition, the National Rifle Association, Christian Coalition, American Israel Public Affairs Committee, and National Right to Life Committee) have little influence. We also know from Gilens's (2012) book that mass-based groups have been largely responsible for the fact that social welfare policy seems to reflect the preferences of low- and median-income citizens more strongly than does government policy in other areas.³²

The tests in that book may provide a better way forward on the question of influence. For instance, some analyses in the book examine not "disagreement" between groups but rather issues on which different income groups diverge in their imputed preference level by more than 10%. I only caution that divergence does not imply low correlation. In addition, frequent overlap between survey questions in the dataset may be a problem for large-*N* analysis. Some issues each generate several similar observations because there are multiple surveys about them. Repeated observations include but are not limited to at least nine questions about NAFTA which all appear as separate victories for the wealthy despite being based on the same policy outcome.³³ For further illustration, I provide code that simplifies the presentation of observations for which elites and median-income citizens diverge by 10 points or more.³⁴ It seems difficult to make normative judgments about the policy process without paying attention to which policies median-income citizens supported or opposed more strongly than elites during the time period in the study.

In closing, "Testing Theories of American Politics" is best where it emphasizes the tentative and imperfect nature of its analysis and where it motivates others to explore

further the question of who really governs. I am grateful for the clarifications that one of the authors was willing to provide. Yet, given existing evidence, average Americans should not believe that it is hopeless to confront or redress through political participation those unfair advantages that elites and organized groups surely do enjoy.

Acknowledgement

The author thanks Yuki Shiraito, Phil Arena, Darren Lim, Saurabh Pant, Kabir Khanna, an editor, two anonymous reviewers, and others for helpful feedback. He also thanks the editor of *Perspectives on Politics* for directing him to a journal that publishes replication studies.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Notes

1. See, for example, headlines such as “Study: US is an oligarchy, not a democracy” and “The US is an oligarchy, study concludes” from *BBC News* and *The Telegraph*, respectively.
2. In lieu of a prohibitively long list of press mentions, I note that *Vox* described it as “The new study about oligarchy that’s blowing up the internet.”
3. Gilens and Page (2014, p. 572).
4. Gilens and Page (2014, p. 573).
5. Gilens and Page (2014, pp. 575–576).
6. Gilens and Page (2014, pp. 576–577).
7. The dependent variable in the study is a dichotomous measure of policy outcome: 1 if a complete change occurs on a given issue within 4 years of an observed poll, and 0 otherwise.
8. Gilens and Page (2014, p. 572).
9. The authors use linear regression because the procedure they use to correct for survey error, described below, cannot be applied to logistic regressions.
10. See Achen (1985), cited in the study, which argues that correlated error and not multicollinearity leads to sign flips with statistical significance preserved.
11. The replication dataset posted by the authors does not facilitate reproduction of the error estimation process, but the study’s appendix describes the approach taken. See also Gilens (2012, p. 253).
12. From the original study: “Data/code for replicating results” available at <http://dx.doi.org/10.1017/S1537592714001595>.
13. See pp. 83–85, available at ftp://public.dhe.ibm.com/software/analytics/spss/documentation/amos/21.0/en/Manuals/IBM_SPSS_Amos_Users_Guide.pdf. These pages show how a measurement submodel can be added around a linear regression submodel to account for measurement error.
14. See rep1.R for the annotated simulation code.
15. The authors’ standard for “essentially zero” is below 0.05. Gilens and Page (2014, p. 575).
16. I thank a reviewer for flagging the number of cases.
17. It also suggests that the linear model used in Gilens and Page may be misspecified. This, however, is not the primary issue with the statistical approach used in the study, as this section has explained.
18. Gilens and Page (2014, p. 576).
19. Gilens and Page (2014, p. 573).

20. See rep2.R for code that replicates these tests.
21. Gilens and Page (2014, p. 572).
22. Gilens and Page (2014, p. 577).
23. This is assuming for simplicity that 0.03 and 0.76 represent regression coefficients in a linear model which does not account for measurement error; the primary analysis above was careful to account for that error.
24. Gilens and Page (2014, pp. 572–573).
25. Gilens and Page (2014, p. 576).
26. Gilens and Page (2014, p. 576).
27. Krugman (2014).
28. The mean preference difference in these cases is over 10%, a preference gap that Gilens considers to be large in his 2012 book.
29. See replication code rep3.R.
30. See, for instance, Gilens (2012), Bartels (2008), or the other works cited in the study.
31. Gilens and Page (2014, p. 569).
32. Gilens (2012, p. 121).
33. Some overlap is less direct: median-income citizens often oppose any form of foreign assistance much more than elites, but aid happens to be a valuable tool of foreign policy (and a much smaller proportion of the federal budget than most Americans believe) so it is used often; many observations in the data involve questions of foreign assistance. In each of the four main policy areas in the study, there are themes which recur across observations. A number of entries in the religious/moral policy category, for instance, involve similar questions about reproductive rights for which outcomes are not independent.
34. See rep4.R.

Supplementary material

The Online Appendix A is available at: <http://rap.sagepub.com/content/srap/suppl/2015/10/05/2.4.2053168015608896.DC1/appendixA.pdf>

The Online Appendix B is available at: <http://rap.sagepub.com/content/srap/suppl/2015/10/05/2.4.2053168015608896.DC1/appendixB.pdf>

The replication files are available at: <https://dataverse.harvard.edu/dataverse/researchandpolitics>

References

- Achen CH (1985) Proxy variables and incorrect signs on regression coefficients. *Political Methodology* 11: 299–317.
- Arbuckle JL (2012) *IBM SPSS Amos 21 User’s Guide*. Armonk, NY: IBM. Available at: ftp://public.dhe.ibm.com/software/analytics/spss/documentation/amos/21.0/en/Manuals/IBM_SPSS_Amos_Users_Guide.pdf (accessed 8 February 2015).
- Bartels LM (2008) *Unequal Democracy: The Political Economy of the New Gilded Age*. New York: Russell Sage Foundation and Princeton University Press.
- Gilens M (2012) *Affluence and Influence: Economic Inequality and Political Power in America*. New York: Russell Sage Foundation and Princeton University Press.
- Gilens M and Page BI (2014) Testing theories of American politics: elites, interest groups, and average citizens. *Perspectives on Politics* 12: 564–581.
- Krugman P (2014) Class, oligarchy, and the limits of cynicism. *New York Times*, 21 April. Available at: <http://krugman.blogs.nytimes.com/2014/04/21/class-oligarchy-and-the-limits-of-cynicism/> (accessed 1 October 2014).