Remarks on Groseclose's *Left Turn*

Justin H. Gross *University of North Carolina at Chapel Hill*
Cosma Rohilla Shalizi *Carnegie Mellon University*
Andrew Gelman *Columbia University*

Begun 6 August 2011, last updated 29 December 2011

In his recent book *Left Turn*, Tim Groseclose concludes that, in a world without media bias, the average American voter would be positioned at around 25 on a 0–100 scale, where 0 is a right-wing Republican and 100 is a left-wing Democrat. In this world, a balanced media might include some TV networks promoting the view that abortion should be illegal under all circumstances and subject to criminal penalties, whereas others might merely hold that *Roe v. Wade* is unconstitutional; some media outlets might support outright discrimination against gays whereas others might be neutral on civil unions but oppose gay marriage; and on general politics there might be some newspapers that endorse hard-right Republican candidates (0 on Groseclose's 0–100 scale) whereas those on the left would endorse positions near those currently held by Senator Olympia Snowe. But instead of this, Groseclose must endure a world where he estimates the average voter is around 50, with all that follows from this, and he attributes this difference to media bias.

Groseclose sets up this stirring climax by formulating and estimating three models. The first model, from Groseclose and Milyo [2005], is an ideal-point model that puts media organizations and members of Congress on that 100 point scale, based on how often they refer to various research and advocacy organizations[1]. The second infers the political positions of voters in different districts based on how many of them voted for Obama, and the estimated positions of their members of Congress. The third model, new to *Left Turn*, is a causal model of how media influence the positions of voters. Groseclose's claims about what our country would look like if it weren't for media bias thus rest on a multi-stage analysis. (Figure 1 depicts his estimation strategy.) He estimates latent quantities, such as how Americans would vote if their views were not distorted by the media, in terms of other latent, estimated quantities, such as the political location of media organizations.

[1] We think "advocacy organization" is a more accurate label than "think tank" for organizations like the American Civil Liberties Union, National Rifle Association, etc.

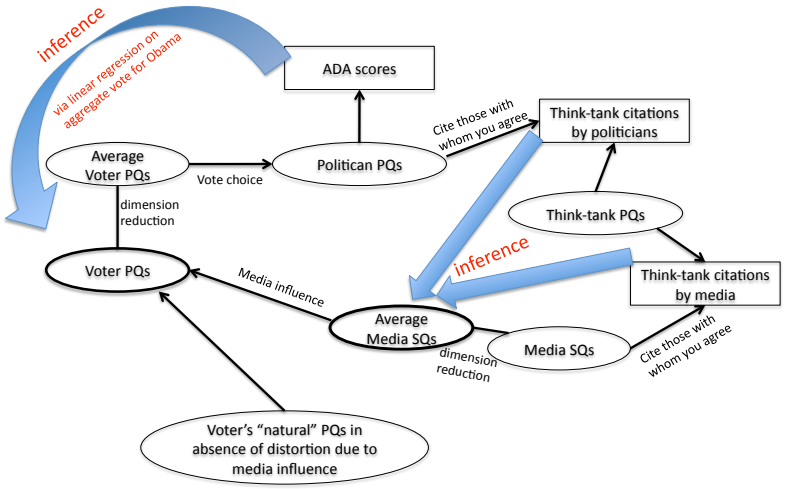

Figure 1: Groseclose's estimation strategy. Variables in boxes are observable; variables in ellipses are latent, and estimated. (The causal network implied by Groseclose's models is subtly different.) PQ and SQ stand for "political quotient" and "slant quotient", respectively.

1 What Is Media Bias?

Groseclose relies heavily (but perhaps unwittingly) on an equivocation about "bias": is the media biased when it *inaccurately represents the world*, or when it *fails to reflect the audience's (natural) opinions*?[2] These are, we might say, respectively direct and indirect sorts of bias, and they differ substantially.

The title, marketing and much of the content of the book invoke the former notion, and the complaint that mainstream media are biased, in this way, against right-of-center political actors and causes has been taken as self-evident among politically-engaged American conservatives since Efron [1971] at least. The accusation is that the mainstream media, while claiming neutrality, are really very liberal in personal ideology *and* in presentation of the news.

Such bias could (in principle) be measured in direct, though labor-intensive, ways: Counting how many Democrats and Republicans, or liberals and conservatives, appear on op-ed pages or TV interviews; tabulating the frequency of political sales terms such as "death tax" or "Operation Iraqi Freedom" or "Affordable Care Act," as compared to more descriptive terms, "estate tax" or "Iraq war" or "the Obama health care plan"; measuring the prominence of positive or negative economic stories (is the bad news always on page A1 and the good news on page D14? does it depend on which party is in the White House?)[3]. All of these would directly measure actions that directly send out

[2]For the purpose of this discussion, we follow Groseclose in hypothesizing the existence of natural opinions that people would hold if they were exposed only to unbiased media.

[3]See for instance Niven [2002], which includes comparison of unemployment coverage for

messages which could influence people.

But *Left Turn*, like Groseclose and Milyo [2005], uses a highly indirect measure of bias, citations to research and advocacy organizations. The operational meaning of "bias" here is that media are biased not if their content provides a systematically distorted estimate of the *news*, but a systematically distorted estimate of the *audience's opinions*. Groseclose, in fact, defines "absolute" bias as the distance between media slant and the "center of American political views" — the latter being found by yet another estimation. In practice, for Groseclose media bias is not a failure of *objectivity*, but of *intersubjective agreement*.

Bias, in Groseclose's operational sense, need not imply bias in the usual sense; the matter is analogous to the difference between statistical lack of diversity in a firm's members and and active discrimination. Just as it takes more than showing that the employees do not reflect the population at large (in terms of age, gender, race, etc.) to prove a practice of discrimination, a lack of representativeness in the media does not imply that they distort the news.

None of this is to say that Groseclose's claims are definitely wrong, just that any interpretation of them is inherently much more fragile than direct measures of bias in communication. One reason the criticisms of Gasper [2011] are so valuable is that they go directly to the data rather than simply assuming that a concordance of references constitutes a political slant.

2 A World in Which Only Media Can Be Biased

In Groseclose's model, all the bias comes from the mass media, not from other political actors. Assume, for argument's sake, that one may meaningfully contemplate the distribution of voters' ideologies in isolation. For Groseclose, this means considering what our collective views would be in the absence of left-leaning media. This holds the media to different standards than other parts of the measurement model. Consider legislators and advocacy organizations, both crucial to Groseclose's estimates of bias. Voters' aggregated ideologies are proxied by the positions of their elected representatives and senators. Legislators' votes on a handful of bills locate them in a latent (one-dimensional) ideological space. Citations of advocacy organizations by these legislators, in turn, place the organizations within that space. Finally, media sources are situated in the space based on their own citations to the organizations. (See Figure 1.) Even if there were no concerns about the propagation of measurement error across these various steps, a crucial flaw in Groseclose's logic is that legislators and advocacy organization are not neutral; their distributions may have just the sort of "bias" he attributes to the mass media.

Again, we are not asserting that the distribution of legislators and advocacy organizations *is* unrepresentative[4]. But Groseclose assumes that these distribu-

presidents and governors of opposite parties, and crime rate coverage for governors.

[4]Though it is easy to come up with accounts of why they might be. Arguably, much of the proliferation of think tanks occurred in direct reaction to the perception among conservatives that academia, as well as the media, had been hostile to their views [Abelson, 2002]. Analo-

tions *are* representative, so that (for instance) slightly left-of-center politicians aren't driven to cite centrist advocacy organizations as the least repellent of an unpalatable lot. This is related to a further striking feature of Groseclose's models, which is that they have only four kinds of actors: voters, media outlets, advocacy organizations, and politicians. Every other sort of institution or organization which might influence political views—schools, the clergy, the public relations industry, to name just three—has no explicit role in the models. They are thus presumed to be either ineffective, or to have no *systematic* effects and so be absorbed into the error terms. These are very strong modeling assumptions, and the evidence in their favor is unspecified.

3 Media Bias as Part of the Political Process

Stepping away from the details of Groseclose's argument, it would make a kind of sense if the news media generally fell on the left side of the political spectrum. Whether this is good or bad is another question, but newspapers generally seem to position themselves on the side of the underdog. Consider, for example, the slogan that a newspaper should "comfort the afflicted and afflict the comfortable." Another way to think about media bias is to think about the communication industry as a whole. We might imagine that "journalism," taken as a whole, leans left, while "public relations," taken as a whole, leans right. And if most of the news media in a country moved from journalism to public relations (as in Berlusconi's Italy), we imagine it could make a difference in the country's politics.

The question, "What would public opinion be like if journalists expressed views comparable to the average American?" while interesting, could perhaps be combined with similar questions such as, What would U.S. politics be like if there were no public relations industry? Or, What would U.S. politics be like if campaign contributions were given equally to the left and the right? Or, What would U.S. politics be like if religious and military leaders were, on average, in the political center?

Our point is: Thinking of all the different institutions affecting political attitudes, it makes a kind of *a priori* sense to suppose that journalism in particular is on the left, while religion (for instance) has been on the right. There are lots of exceptions, from Martin Luther King on one side to Silvio Berlusconi on the other, but our first guess would be that journalism is one of the left-leaning institutions in the U.S. And it makes sense for Groseclose, as a conservative media analyst, to want to shift journalism to the right, just as, from the other direction, a liberal businessman might want to persuade businesses to move in the other direction.

Policy and politics are multidimensional. For example, *Slate* magazine notoriously polled its staff before the 2008 election and found 55 out of 57 supporting Obama. On the other hand, a *Slate* writer (who we would guess is an Obama

gously, the distribution of legislative candidates may filtered by the need to finance campaigns through large donations.

supporter) wrote, "If we can find other ways of overcoming the simmering resentment that naturally accompanies wage cuts, workers themselves will be better for it in the long run" [Fisman, 2011]. The "we" at the beginning of the sentence does not include the "workers" at the end of the sentence. This is just an anecdote ($n = 1$, and not a randomly sampled $n = 1$ at that) but it does reflects a general attitude in the big media, which by default take the perspective of the employer or rich person rather than the employee or poorer person. (Think of the famously obnoxious lifestyle pieces in the *New York Times*.) The multidimensionality of political attitudes should not discourage us from studying bias, but does complicate it.

In the U.S. context there is arguably an asymmetry in political bias, with reporters who are Democrats—consistent with Groseclose's claims, a survey a few years ago found that twice as many journalists identify as Democrats than as Republicans—biasing their reporting by choosing which topics to focus on, and Republican news organizations (notably Fox and other Murdoch organizations) biasing in the other direction by flat-out attacks.

We have never been clear on which sort of bias is more effective. On one hand, Fox can create a media buzz out of nothing at all; on the other hand, perhaps there's something more insidious about objective news organizations indirectly creating bias by their choice of what to report. But we have long thought that this asymmetry should inform how media bias is studied.

4 Counterfactuals

The new scholarly contribution of the book, over the QJE paper, is the model of media influence in ch. 20 (lead up to by chs. 18–19, and used thereafter). This is a causal model, used to calculate many counterfactuals. Model-based counterfactual inference is of course a legitimate part of science [Morgan and Winship, 2007], but one must ask, *why this particular model?*, and to this the book has no good answer.

Groseclose's model is that the average opinion of a population is a convex combination of the average of their natural opinions, and the average of the media they are exposed to; the relative weight of the media is λ. We have not found any argument by Groseclose for the convex-combination model, in the book or elsewhere. A static model of influence or persuasion is already a great simplification of a dynamic process. Substantively, the phenomenon of group polarization [Sunstein, 2002, 2009] shows that discussion among like-minded individuals often results in the average view of the group becoming *more* extreme—sometimes more extreme than the initial views of any member. This finding is robust and has been demonstrated in both a wide range of experiments and in high-stakes situations like judicial panels [Sunstein, 2009]. This is incompatible with any convex-combination model of influence, so Groseclose must assume that media influence is somehow distinctly different from interpersonal influence, though this makes his use of signaling-game findings odd.

The convex-combination model for average opinions must be an aggregation

of individuals' beliefs. If the individual-level process is also a convex combination model, then very serious issues of aggregation arise. (See the online appendix for details on what follows.) For the aggregated and individual models to coincide, voters must not be influenced by politicians or by each other, or chose what media they consume, or trust or attend to some media sources more than others. Violating the influence restrictions, in particular, can render the whole model unidentifiable. The convex-combination model could of course hold as an emergent property of the group, not displayed by its members, but then Groseclose could not use of individual-level experiments (like the signaling game) to estimate its parameters.

Groseclose wishes to evaluate the influence of the national media on national elections. His data sources for the model do not address this. He thus needs λ to be at least roughly "projectible" across many contexts, say the U.S. in the past two decades. He goes beyond this to posit that λ is a universal constant, comparable to the constant in Newton's law of gravitation. He does not share why he thinks this. It would not follow from aggregation without special assumptions, but it would be even more mysterious as an emergent property.

The fundamental problem with the use of the convex-combination model is that Groseclose makes no effort to *check* his model of media effects, i.e., to systematically compare it to data and show that departures from its predictions are unsystematic and effectively noise. Once again, it's fine to use models to infer counterfactuals, but then one needs to look into possible problems.

Groseclose does compile estimates of λ implied by different data sources. Under his assumptions, these should all be compatible; he does not try to check this, even roughly by finding standard errors. Indeed, the experiment by Gerber et al., which Groseclose praises as exemplary, yields a nonsensical estimate of λ, greater than 1. Groseclose attributes this to sampling noise, but does not back this up by calculating any measure of uncertainty. It might be more useful to consider an out-of-bounds estimate as a warning of a problem with the model.

5 Reception of the Book

Political observers of all ideologies recognize the importance of news media in influencing public opinion. One reason we believe that the Groseclose and Milyo work has been taken so seriously (with over 200 citations so far in Google Scholar) is that it was a relatively sophisticated quantitative take on an important problem.

Consider, for example, the blurb by economist Steven Levitt (who has collaborated with Groseclose in the past): "[*Left Turn*] is not, however, a right-wing rant by any means. Rather, it is a carefully researched and amusingly written book by a highly regarded academic. ... liberals will not like what Groseclose has to say, but that is all the more reason why liberals should read his book." Consider what is left unsaid: Levitt does not write that the book's claims are scientifically proven or even substantively correct, nor does he endorse Groseclose's main conclusion, that, without media bias, the average American would

think and vote like Fox News personality Bill O'Reilly. Rather, Levitt writes that the book is "carefully researched," which praises the process and not (necessarily) the product. Similarly, on a blog discussion of the Groseclose and Milyo article [Gelman, 2007], one of us expressed some concern about the methodology but wrote, "I'm a bit more positive than [Nyhan] is about the paper, I think because the problem of studying media bias is tough, and I'm impressed about what Groseclose and Milyo did manage to do."

Academic communities sometimes hold our writings to high levels of rigor (internal validity) but a low level of correspondence to reality (external validity). Sometimes this is even appropriate, such as speculatively exploring the implications of hypotheses which might later be refined into explanatory theories; on such a basis it seemed reasonable to vaguely endorse Groseclose and Milyo's efforts as published in 2005 without having to believe its conclusions. Emphasizing rigor over accuracy becomes a failure-mode of scholarship, however, when internal validity is *mistaken* for external validity, and a rigorously worked-out conjecture for the correct explanation (cf. Rogeberg and Melberg 2011). Misconceptions which arise in this way can be incredibly tenacious, precisely because of their scholarly form [Boudon, 1986/1989, Hamilton, 1996]. In *Left Turn*, Groseclose makes the commendable effort to connect the dots and think about the political implications of his work. That's fine, but then we find ourselves less inclined to be charitable about flaws in the research. "Carefully researched and amusingly written" can be enough for an academic article, or for a popular book, but it does not suffice for a popular book that purports to provide scientific proof of a controversial claim[5].

6 Conclusion

In moving from an academic journal to a popular work, Groseclose has added a partisan, even rabble-rousing, dimension to his work. This political activism is appropriate given the importance of Groseclose's claims: if the biased media are indeed moving the average American from the position of Bill O'Reilly to that of Joe Lieberman, this is a big deal. There is a long tradition of left-of-center social scientists (recently, for instance, William Julius Wilson and Paul Krugman) expressing strong recommendations along with their research findings, and Groseclose can certainly do likewise from the other side. We have no criticism of Groseclose's care in describing his own political positions[6].

[5] We agree with Levitt that *Left Turn* is well-written; unfortunately the cleanest parts of the writing correspond to the simplest, most speculative, parts of the model (for example, Groseclose's assertion that individuals' political attitudes are a weighted average between media positions and voters' underlying true ideologies), while the more complicated, data-based parts of the analysis are described less clearly. We refer readers to Gasper [2011], who is admirably clear on the details of the data behind the analysis.

[6] However, we think Groseclose went over the line as a partisan advocate when, on a Fox News interview, he said that the left-leaning organization Media Matters is "violating the spirit of the law" in doing political advocacy while being classified as a tax-exempt group under IRS Code Section 501(c)(3). Apparently, though, this same tax exemption is claimed by many conservative educational groups as well, including the Heritage Foundation, Cato

In addition to its crowd-pleasing rhetoric, *Left Turn* also aims to be a serious work of political science. In particular, it aspires not just to be a theoretical exploration of the consequences of certain hypotheses ("what would happen if the media were really, really biased?"), but an accurate explanation of why, in fact, America has the politics it does. As noted above, the book relies heavily on a set of theoretical models that are multiple steps removed from actual media bias. Ultimately the work must be judged not (just) by whether its proposals are normatively compelling, or by whether its hypothetical speculations are internally coherent, but on the fit of its models to political reality.

References

Donald E. Abelson. *Do Think Tanks Matter? Assessing the Impact of Public Policy Institutes*. McGill-Queen's University Press, Montreal, 2002.

Raymond Boudon. *The Analysis of Ideology*. Polity Press, Cambridge, England, 1986/1989. Translated by Malcolm Slater from *L'idéologie, ou, L'origine des idées reçues*, Paris: Fayard.

Edith Efron. *The News Twisters*. Nash, Los Angeles, 1971.

Ray Fisman. Raises don't make employees work harder. *Slate*, March 2011. URL `http://www.slate.com/articles/business/the_dismal_science/2011/03/raises_dont_make_employees_work_harder.html`.

John T. Gasper. Shifting ideologies? Re-examining media bias. *Quarterly Journal of Political Science*, 6:85–102, 2011. doi: 10.1561/100.00010006. URL `http://www.andrew.cmu.edu/user/gasper/WorkingPapers/IdeologicalShiftv2.pdf`.

Andrew Gelman. Measuring media bias, May 2007. URL `http://andrewgelman.com/2007/05/measuring_media/`.

Tim Groseclose and Jeffrey Milyo. A measure of media bias. *The Quarterly Journal of Economics*, 120:1191–1237, 2005. URL `http://www.sscnet.ucla.edu/polisci/faculty/groseclose/pdfs/MediaBias.pdf`.

Richard F. Hamilton. *The Social Misconstruction of Reality: Validity and Verification in the Scholarly Community*. Yale University Press, New Haven, 1996.

Stephen L. Morgan and Christopher Winship. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press, Cambridge, England, 2007.

Institute, American Family Association, and Family Research Council. If Media Matters is violating the spirit of the law, so are all these other groups. Groseclose also misattributed to Media Matters a headline of a critical article that actually appeared on the website of the Center for American Progress.

David Niven. *Tilt? The Search for Media Bias*. Praeger, Westport, Connecticut, 2002.

Ole Rogeberg and Hans Olav Melberg. Acceptance of unsupported claims about reality: A blind spot in economics. *Journal of Economic Methodology*, 18:29–52, 2011. doi: 10.1080/1350178X.2011.556817.

Cass R. Sunstein. *Going to Extremes: How Like Minds Unite and Divide*. Oxford University Press, Oxford, 2009.

Cass R. Sunstein. The law of group polarization. *The Journal of Political Philosophy*, 10:175–195, 2002. doi: 10.1111/1467-9760.00148. URL `http://ssrn.com/abstract=199668`.

A The Rotten Tomatoes Paradox, and Uncertainty

Frequent visitors to the website, Rotten Tomatoes, may be surprised to discover that a movie with a 15% positive rating is not necessarily better reviewed than than one which gets, say, 45% positive. This is because dichotomizing reviews as either just good or bad, period, discards lots of information. A strong consensus that a film is not great (say, mostly C-) will result in a low percentage of positives. On the other hand, a movie with a lot of terrible reviews but a strong lukewarm contingent (say, 55% Fs and Ds, 45% C+) can have a much higher "positive" percentage.

While the subject-matter is much less important, this example has *exactly* the structure of voters located on a one-dimensional latent political space being forced to chose between two candidates.

More generally, this points to the importance of uncertainty assessments. Even if one were to accept all of Groseclose's modeling assumptions, the reader of *Left Turn* has no idea of what the views of the hypothetical unbiased American voter are. The reader knows Groseclose's *point estimate* (≈ 25, as we said), but this estimate, like every other, is uncertain and surrounded by some margin of error. Since it is a function of earlier estimates, themselves imprecise[7], the obvious and correct thing to do is to propagate the uncertainty. But *Left Turn* gives no sense of how much uncertainty attaches to its claims, even of a purely statistical sort. Equally seriously, Nyhan [2005] and Gasper [2011] have shown that estimates from Groseclose's first, ideal-point model are highly sensitive to details of which data are included in the analysis, adding systematic to statistical uncertainty. We do not know how big the margin of error for Groseclose's estimate of the location of Unbiased America should be, but it surely matters if a good confidence interval is $[22, 28]$ or $[10, 90]$, or indeed $[0, 100]$.

B Aggregation of Convex-Combination Influence Models

Presumably, the model relating the average opinion of groups of voters to their average media exposure must arise somehow through aggregating the relationship between individuals' opinions and their media exposures. (After all, even if Groseclose believed in some sort of Durkheim-style "collective consciousness", what he's *measuring* is an aggregate of individual opinions.) The simplest possibility, which seems to be assumed by the way Groseclose freely mixes group- and individual-level observations, is that a convex combination model applies to individuals[8]. In this appendix, we consider what must be assumed in or-

[7] Ignoring all specification issues, locations estimated from ideal-point models, even NOMINATE, typically have much bigger margins of uncertainty than is often appreciated [Gross, 2010].

[8] We note again the phenomenon of group polarization, in which individuals' opinions, after discussion, shift to values which are more extreme than any that they started with,

der to aggregate individual convex-combinations models into the *same* convex-combination model at the group level.

The first striking limitation is that, in this model, the *only* sources of individuals' observed opinions are their natural opinions and the media. They are not, for instance, influenced by historical events, office holders, or political campaigns. More crucially, individuals are not allowed to influence *each other*. Allowing them to do so, however, lands us directly into the mire of unidentifiability classically mapped out by Manski [1993], and which no amount of subsequent work has managed to noticeably drain [Manski, 2007, Blume et al., 2010]. Specifically, it becomes impossible to distinguish between the influence of a covariate (such as shared media exposure) and social influence within the group.

For the sake of argument, suppose that each individual's opinion really just is a convex combination of their natural opinion and the average opinion of the media they are exposed to (plus, presumably, noise). Further suppose that every individual has the same λ. For the aggregated, population-level model of average opinion to have the same form, with the same λ, and with the *group's* average media exposure as the driving variable, we must assume that individual media exposures are uncorrelated with their individual natural opinions. This in turn entails that people must be unable to select media sources because they are congenial to them, which rather contradicts the whole estimation procedure for politicians in the first half of the book. Similarly, it is implicit that people cannot assign *different weights* to different media sources depending on how much they trust them (depending, e.g., on the gap between the media's PQ and their natural opinions). So to aggregate the individual-level form of the model to the population-level form, which is what Groseclose needs, one must assume (i) that λ is constant across people, (ii) that the *amount* of media exposure is uncorrelated with natural opinions, (iii) that the *effectiveness* of media one is exposed to is unrelated to natural opinions. If these assumptions fail, then at the aggregate level, the average opinion would be a convex combination of the average natural opinion, and *some* weighted average of the media, but the relevant weights would not be the market share.

Constancy of λ Casual empiricism shows that some people are stubborn in their views, some people are credulous, some people give little attention to the media and others are news-junkies, etc. This suggests that even if the model is valid on the individual level, λ should *not* be constant across people. Indeed, given the extensive individual differences in basically every psychological trait involving higher mental faculties, it would be astonishing if λ *were* constant. Thus, even if there is no correlation between natural views and news sources, no weighting by trust, etc., a constant lambda for populations could only be a probabilistic approximation. In fact, constancy would require non-trivial substantive assumptions, decoupling individual λ values from natural views, available media, etc.

---

which cannot be reconciled with a convex-combination process.

C The Errors-in-Variables Issue

Unless Groseclose wishes to assert both that (i) his PQ scores for media outlets are *really* the causally-relevant variable for influencing audiences, and (ii) he has measured PQ scores without error, his average-media-opinion is only a proxy for the media's true opinion (or whatever the exact causal variable is). Similarly, to the extent he measures audiences' opinions by votes, there is another level of proxying. So he really has an error-in-variables model, or even a structural-equations/latent-variables model, and would need to estimate accordingly, which he does not. In many cases, the relations between latent variables are *stronger* in such situations than the relations between their noisly-measured observable proxies, so it is even possible that Groseclose's procedure underestimates media effects in his causal model. It is, however, certain that what he does ignores the difference between his proxies and his real variables.

Supplementary References

Lawrence E. Blume, William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides. Identification of social interactions. In Jess Benhabib, Alberto Bisin, and Matthew Jackson, editors, *Handbook of Social Economics*, page forthcoming. Elsevier, Amsterdam, 2010. URL `http://ssrn.com/abstract=1660002`.

John T. Gasper. Shifting ideologies? Re-examining media bias. *Quarterly Journal of Political Science*, 6:85–102, 2011. doi: 10.1561/100.00010006. URL `http://www.andrew.cmu.edu/user/gasper/WorkingPapers/IdeologicalShiftv2.pdf`.

Justin H. Gross. *Cues and Heuristics on Capitol Hill: Relational Decision-Making in the United States Senate*. PhD thesis, Carnegie Mellon University, 2010.

Charles F. Manski. *Identification for Prediction and Decision*. Harvard University Press, Cambridge, Massachusetts, 2007.

Charles F. Manski. Identification of endogeneous social effects: The reflection problem. *Review of Economic Studies*, 60:531–542, 1993. URL `http://www.jstor.org/pss/2298123`.

Brendan Nyhan. The problems with the Groseclose/Milyo study of media bias. Weblog post, 2005. URL `http://www.brendan-nyhan.com/blog/2005/12/the_problems_wi.html`.