POLITICAL OUTCOME BIAS IN GRADING:

IDENTIFYING PROBLEMS, SUGGESTING SOLUTIONS

Mark Carl Rom, Georgetown University
Mailing address:
37th and O Street, NW
312 Old North Hall
Washington, DC 20057
Email: romm@georgetown.edu
Phone:202-213-8767

Paul Musgrave, Georgetown University
Mailing address:
37th and O Street, NW
Department of Government, ICC 681
Washington, DC 20057
Email: Paul.musgrave@gmail.com

Author information:

Mark Carl Rom is Associate Professor of Government and Public Policy at Georgetown University. He has published numerous books, chapters, and articles on American social welfare policy, health, and finance. His most recent publications are "Grading More Accurately," *JPSE*, 2011, 7: 208-223 and "The Scholarly Conference: Do We Want Tradition and Authority or Markets and Democracy?" *JPSE*, forthcoming.

Paul Musgrave is a Ph.D. candidate in the Department of Government at Georgetown University. A graduate of Indiana University, he studied at University College Dublin as a George Mitchell Scholar. A former editor at Foreign Affairs, he was special assistant to the director at the Richard Nixon Presidential Library and Museum before beginning his doctoral studies.

POLITICAL 'OUTCOME' BIAS IN GRADING:

IDENTIFYING PROBLEMS, SUGGESTING SOLUTIONS

Abstract:

Political bias in the academy is a topic of great controversy. Many conservatives have argued that liberals dominate American campuses and use their classrooms to indoctrinate students or to discriminate against those with differing political beliefs. Liberals have responded by calling studies that purport to demonstrate these claims as flawed or as attacks on academic freedom. Regardless of the magnitude of campus political bias, it is ill-advised for the scholarly community to argue that it is immune from bias because scholars simply are fair. This paper focuses on one element of political bias: partisan 'outcome' bias in grading. We proceed in several steps. First, we provide an overview of the problem of grading bias and the concerns about political bias. Next, we consider the attributes of political grading bias and the forms that it can take. We present the analytics of partisan outcome bias and logic of using 'paired' assignments to eliminate outcome bias. We provide data and analysis concerning outcome bias from a course [Author] has taught. We conclude with a discussion of the ethics of grading.

POLITICAL OUTCOME BIAS IN GRADING:

IDENTIFYING PROBLEMS, SUGGESTING SOLUTIONS[1]

INTRODUCTION

"Mark clearly has a liberal bias. I am a Christian, moderate Republican and think that everyone's personal views, including mine, influence my thoughts on ethical issues. Mark graded me down because I acknowledged that my personal beliefs influence my policy decisions" (Anonymous 2006).

"Students will be graded solely on the basis of their reasoned answers and appropriate knowledge of the subjects and disciplines they study, not on the basis of their political or religious beliefs" (Students for Academic Freedom, n.d.).

[Author's] initial reaction to the anonymous posting on RateMyProfessors.com was dismay and denial, as he firmly believed that his grading was neutral, and could not understand how one of his students thought that he was biased. He was tempted to conclude that this student was merely disgruntled. On further reflection, the posting raised two questions. What if the student actually were right? And even if the student were wrong, why might students perceive that their professor was politically biased?

Political bias in the academy is a topic of great controversy (for a review, see Mariani and Hewitt 2008; see also McCardle 2011a, McCardle 2011b, Tierney 2011 and the readers' responses to these blogs; for other sources and commentary see Haidt 2012). Prominent conservatives have argued that liberals dominate American campuses and use their classrooms to indoctrinate students or to discriminate against conservatives (see especially Horowitz 2007). To the extent

2

they have deigned to reply, liberals have responded by suggesting that studies purporting to demonstrate these claims are overblown, shoddy, or attacks against academic freedom (Lee 2006; Vitiello 2007).

Still, a few experimental studies have found evidence of partisan bias in terms of college admissions (Munro, Lasane and Leary 2010) as well as discrimination against conservative ideas in IRB recommendations (Ceci, Peters and Plotkin 1985) and the publication of research (Abramowitz, Gomes and Abramowitz 1975). More recently, Inbar and Lammers (2012) found that "In decisions ranging from paper reviews to hiring, many social and personality psychologists admit that they would discriminate against openly conservative colleagues. The more liberal respondents are, the more willing they are to discriminate" (Inbar and Lammers 2012)

Regardless of the magnitude of campus political bias, it is ill-advised for the scholarly community to argue (or insist!) that it is immune from bias because scholars simply are fair. Certainly, political scientists and the broader public would react with incredulity if government officials with virtually unfettered discretion responded to charges of favoritism or bias by merely stating that they were above reproach. Better, certainly, that professors take seriously the potential for political bias in grading, and take steps to prevent—or at least ameliorate -- it. This paper provides theory and evidence about how instructors can do so, especially by adopting a simple protocol: using assignments that require students to take opposing sides on the same topic.

This paper proceeds as follows. First, we provide an overview of the problem of grading bias and the concerns about political bias. Next, we consider the attributes of political grading bias and the forms that it can take. We present the analytics of partisan outcome bias and logic of using

3

'paired' assignments to eliminate outcome bias. We provide data and analysis concerning outcome bias from a course [Author] has taught. We conclude with a discussion of the ethics of grading.

THE PROBLEM OF BIAS IN GRADING

We assume that most professors would deny that their grading is politically biased, although we have found no published research on this point. Other forms of grading bias have been demonstrated in numerous ways, as we shall see below, and other studies have identified that bias can exist even among persons who believe they are unbiased. Moreover, given that a large majority of professors in the social sciences and humanities are liberal and that a majority are Democrats, it is worth examining whether political bias may taint those professors' grading schemes. Further, even if bias is minimal or nonexistent, a substantial proportion of students and the general public believes that professors are politically biased, and therefore taking steps to reduce this perception is worthwhile.

Definition of Grading Bias

Grades are estimates of student achievement, an unknown parameter. In statistical terms, an estimator is the function used to produce the estimates of this parameter. Good estimators are in general efficient, unbiased, and consistent (Kennedy 2001, 11-20). In statistics, bias is a theoretical construct: bias exists if the mathematical properties of the estimator indicate that the resulting estimates systematically (i.e., not randomly) deviate from the true (i.e., known) parameter. An estimator is unbiased if it yields estimates that, on average, do not deviate from the parameter (that is, the expected value of the estimate equals the parameter); in contrast, an estimator is biased if it produces estimates that on average are either higher or lower than the true

4

parameter. By analogy, grading bias results from a systematic failure to estimate the true nature of student achievement based on available evidence.

Grading bias is the result of incorporating illegitimate factors into an instructor's assessment of students' work. The immediate consequence of grading bias is that grades do not reflect the student's true performance in mastering course material. The longer-term consequences may include forcing students to choose between doing their best work or playing to their instructor's biases, the balkanization of university life as students self-select out of courses and disciplines they perceive to be unwelcoming, and the weakening of the academy's claim to have an impartial commitment to dispassionate inquiry.

The key term in the definition of grading bias is 'illegitimate.' For grading to be biased, the instructor must be unable to justify the factors that enter into assessing students' work. It is not 'grading bias' for an instructor to announce in a graduate-level course on data analysis that final papers must employ Bayesian principles to receive full marks, nor would it be illegitimate for a professor of a course on political institutions to require students to prepare an assignment using game-theoretic tools. Such standards are defensible because they allow for instructors to assess whether students have learned the material covered in a course and they allow students to know the standards by which they are assessed. However, failing to announce such a policy and then marking down frequentist or historical-institutionalist approaches would be illegitimate, since students would have no reasonable or clearly stated expectation that such a standard would be employed.

Bias can be difficult to prove. However, under certain circumstances, bias can be demonstrated analytically and empirically. For example, consider students who come in two types—say, X and Y. Let us assume that the X students receive generally higher scores from the professor than the

5

Y students. Does this indicate bias? Not necessarily. Perhaps the X students actually deserve higher scores, in which case the estimator used to award scores is unbiased. To determine whether bias exists, let us assume that the X and Y characteristics can be hidden from the professor. Under this scenario, if the grades do *not* show differences between the X and Y students, we can conclude that it was the knowledge of the attribute that biased the professor's grades; if the grades of X and Y are still different even though the student types are concealed, we should be more confident that the grades truly reflect student performance.

Sources of Non-Political Grading Bias and Potential Remedies

A substantial literature has examined the potential sources and actual occurrences of grading bias. There are several distinct non-political mechanisms of bias, with each requiring a different remedy.

The first -- and the most heavily studied -- set of factors that can lead to biased grading involve *inherent attributes of the student*: race, ethnicity, nationality, gender, and physical attractiveness (for a summary of the literature, see Fleming 1999: 85-86). These problems are no less worrisome for their familiarity: for instance, Spear (1984) found that male and female teachers both tended to perceive more highly the work of male students studying in more typically 'masculine' subjects (such as physics), while female students were more highly rated in 'feminine' subjects. Preconceptions of the quality of the student can also bias grading. Here, *halo* (e.g., "Student X is terrific!") or *horn* (e.g., "Student Y just doesn't have 'it'.") effects lead the professor to give higher or lower grades than the student's work merits. One study concludes that 'decades of halo-effects…provide substantial evidence for a pervasive effect of prior knowledge of a target on judgments of the targets' subsequent performance' (Archer and McCarthy 1988, as quoted in Fleming 1999: 86).[2]

6

There are no perfect remedies for bias involving the inherent attributes of the student. The best approach is – whenever possible – to conceal these attributes from the grader by having written assignments submitted by student ID code rather than by name. This allows the grader to assess the work before linking the ID number to the student's name.[3] It is not possible to conceal the identity of the student, of course, for oral presentations, for assignments that require close collaboration of the student with the instructor, or for classes that are so small that student identities might be inferred from the assignment itself. Anonymous submissions are beneficial, but they are not always possible.

Contrast effect bias results when an assignment is scored higher (or lower) than deserved due to where it fits in the grading sequence. For example, research has indicated that 'average' papers are scored lower when they are preceded by several high-quality papers and higher when graded after a series of low-quality papers (Hales and Tokar 1975; Hughes *et al.* 1980; Daly and Dickson-Markman 1982). As Helson describes it, when individuals assess items that are "significantly different from the established norm [they] adjust or contrast the new stimulus with a more extreme position than is warranted by the object's true value" (Helson 1964, quoted by Fleming 1999: 87). As a result, the score a paper receives is influenced not just by its quality, but by the sequence in which it is assessed.

Once again, no perfect remedy is possible for any given assignment: inevitably, worse papers will follow better ones, and vice versa. Over the course of the semester, however, the overall impact of contrast bias will be mitigated if each assignment is graded in random order, so that the probability that any given paper will is graded after another superior or inferior paper is minimized.

7

Finally, *presentation bias* takes place when the appearance of the submitted work (e.g., length, neatness, quality of formatting and printing, and so forth) influences grading. The solution in this case is for the instructor to standardize, to the extent possible, the format, structure and length of each assignment. Requiring assignments to be typed, rather than hand written, will also reduce presentation bias.

Political Grading Bias

Critics of the academy have focused their concerns on partisan or ideological bias, with the main implication that liberal, Democratic professors tend to discriminate against conservative views and Republican students. We accordingly define two primary types of political bias:

In *outcome* bias, professors discriminate among students based on the conclusions of their arguments. Ideological or partisan professors provide higher scores than deserved to students' *work* that agrees with the professors' ideological or partisan preferences, and lower scores to the work that disagrees with these preferences.[4] For example, outcome bias would exist if a Republican professor rewarded essays concluding that voting ID standards should be tightened and enforced to prevent voter fraud, or if a Democratic professor favored essays concluding that such voter ID policies were a discriminatory effort to suppress voting among disadvantaged groups.[5]

In *affinity* (or *caricature*) bias, professors discriminate among *students* based on the students' affiliation with the professor's partisan identification, so that professors award scores higher than deserved to students of matching identifications, and lower scores to students who have a mismatching partisanship. For example, if a Democratic student and a Republican student each wrote two essays, with one favoring tighter voter ID requirements and the other opposing them, affinity bias implies that the Democratic professor would favor the essays written by the

8

Democratic student over those written by the Republican student irrespective of the specific conclusions reached. Much as Campbell et al (1960) argued about the role of perceptual filters in shaping partisans' view of the world, the affinity bias hypothesis suggest that there may be such a mechanism at work in shaping instructors' assessments of students (and, *a fortiori*, of students' views on the issues they are to write about). This is a specifically political hypothesis about grading behavior, since it implies that students' and instructors' political beliefs and behaviors outside the classroom can have important ramifications inside the instructional environment.

These two distinct types of bias imply that professors can exhibit no bias, either outcome or affinity bias, or both biases. Although other forms of political bias are possible, we focus on outcome bias in this paper.[6]

For all the passion that the topic of political bias in the academy has engendered, virtually no research exists on the matter of partisan bias in grading. A single major study has assessed the relationship between student ideology and grading patterns at a single large public university: it concluded that conservative students received grades that were equal to or higher than their liberal counterparts (Kemmelmeier, Danielson and Baston 2005).[7] This study does not seek to measure bias directly; simply observing differences across students or professors does not necessarily demonstrate bias or the absence of it. Neither students nor professors are randomly assigned to classes, for example, but instead choose to take or teach particular courses for various reasons. If liberals (or conservatives) generally selected easier (or more difficult) classes, their average grades would differ across the categories due to selection, not political, bias. Similarly, if liberals tend to take 'liberal' courses (or courses from liberal professors), while conservatives do the same, then there might be no observable bias in overall grades even though bias within individual courses could be endemic. Finally, if students modify their assignments to

match their perceived ideology of their professors, then any outcome bias would be masked in the assignment of grades: liberal professors would receive liberal papers, and conservative professors conservative ones. Once again, bias would be unobservable directly, but still present and important.

More recent work by Bar and Zussman (2012) finds differences between how Republican and Democratic professors at one elite university assign grades, with Republicans favoring a much more markedly unequal distribution of grades, which the authors analogize to Republican and Democrats' preferences over merit and egalitarianism. This study shows partisan differences, but not partisan bias.

Our experience suggests that faculty often believe that they are tougher on the students who agree with them. If this were true, then outcome bias would actually favor the group the professor politically opposes. Little evidence supports this prospect, given the likelihood of confirmation bias. As Mahoney (1977: 161) defines it, confirmation bias is "the tendency to emphasize and believe experiences that support one's views and to ignore or discredit those that do not. The effects of this tendency have been repeatedly documented in clinical research" (Mahoney 1977: 161). Confirmation bias implies that professors are likely to be more critical of work that contradicts their own views, to be more skeptical of sources and arguments that oppose their beliefs, and, perhaps, to give higher scores to those papers that are in accordance with their political understandings (see for example Johnson 1996; Lundgren and Prislin1998; Jonas *et al*. 2001; Lord, Ross, and Leeper 1979, and the discussion in Kelly-Woessner and Woessner 2006).

If the political preferences of college professors were randomly distributed, and bias was randomly distributed across professors, then bias would average out and any bias in individual courses would have no net impact on the students' overall grade point average. But professorial

10

political preferences are clearly not random. Although the literature on the political preferences of college professors has generated substantial controversy (see for example Mariani and Hewitt 2008; Klein and Western 2004-5; Horowitz and Lehrer 2002; Cohen-Cole and Durlauf 2005: 4; Klein and Stern 2004-5; Rothman, Lichter, and Nevitte 2005; Lee 2006; and Zipp and Fenwick 2006), the best recent research from a large, random sample of faculty indicates that college professors are much more likely to identify with the Democratic party than the Republican party. Thirty-two percent are strong Democrats compared to only ten percent who are strong Republicans, while over 70 percent are at least somewhat aligned with the Democrats compared to about 20 percent who at least partially self-identify as Republicans (Gross and Simmons 2007, 31). Sixty-two percent of faculty considered themselves at least 'slightly liberal' compared to about 20 percent who identified themselves as at least 'slightly conservative,' with the remainder seeing themselves as 'middle of the road' (Gross and Simmons 2007, 26). Among political scientists, 50 percent identified themselves as Democrats and 6 percent as Republicans, with the remainder identifying themselves as independents (Gross and Simmons 2007: 31). Moreover, professors' policy positions tend to be more liberal than their party ID might suggest (Rothman, Kelly-Woessner, and Woessner 2011). So if professors are politically biased in grading, and if Democrats and Republicans as well as liberals and conservatives are equally likely to be biased, the odds are about 3-to-1 that the bias will come from the left side of the aisle.

Political scientists, along with their colleagues in the social sciences and the humanities, should be particularly concerned about political grading bias to the extent that disciplinary consensus over standards (i.e., what is and should be taught) or the content of a discipline is seen as subjective. If we arrange undergraduate courses along these two dimensions (i.e., consensus/dissensus and objectivity/subjectivity), a discipline like chemistry would reflect high

11

levels of both consensus and objectivity, while a course such as a poetry workshop would score much more heavily toward dissensus and subjectivity. It would be easy for a chemistry professor to avoid political bias in grading, especially if standard safeguards (such as anonymous paper submission) were employed. Contrariwise, a professor in an MFA course might well be able to bias grades and to do so with little fear of sanction, since uncovering the causal weight of bias in a highly subjective environment is difficult or impossible. Political science courses vary widely in their scores on this notional two-dimensional index. Some courses, such as one teaching game theory, are more like chemistry (or, perhaps, algebra), while others (e.g., an area studies or an ethnographic course) might move toward the other extreme.

Many assignments in political science classes are likely to have partisan implications that might be especially susceptible to political grading bias, however. That is certainly the case with highly normative work, such as a student writing a paper on the topic "Should school voucher programs be adopted?" But the political bias need not just involve questions of values, as so often even the facts of the matter are also subject to considerable dispute (for example, see Belfield and Levin 2005 for a discussion of the partisan debates over evidence in the voucher controversy). At any rate, discussions over facts are themselves subject to confirmation bias. If the students sincerely display their political views, then the professors' own views possibly are source of grading bias. One of the potentially pernicious consequences of political bias in grading, however, is that the students might modify their own work in an attempt to curry favor with the professor. The potential for this is not trivial, because students apparently are pretty good at 'guessing' their professors' political ideology (Kelly-Woessner and Woessner 2006, 496; Woessner and Kelly Woessner 2009; see also Jaschik 2006) and because students prefer higher grades to lower grades. One survey shows that, indeed, a quarter of college students believe they need to 'parrot'

12

their professors' views in order to get a good grade (ACTA 2006, 2; see also the commentary by Leef 2006 and Balch 2006).

Finally, concerns about bias persist independently of classroom realities. Substantial proportions of the public believe that professors are politically biased. According to a poll by the American Association of University Professors, 49 percent of Republicans and 27 percent of Democrats describe political bias in the classroom as a "very serious problem" (Smallwood 2006). Consequently, even if the reader does not accept that bias is a problem in actuality, the scope of public concern over perceived bias should prompt a thoughtful response.

We have now established that there are several different types of mechanisms through which bias in grading can be manifested, and we have established that there are grounds for further investigation of political bias in grading. In the next section, we consider the analytics of political grading bias and suggest a measure to mitigate outcome bias.

THE ANALYTICS OF POLITICAL GRADING BIAS

In order to understand political grading bias, we first define hypothetical students and their relationships with their professor's partisanship. We denote students (*S*), whether or not the conclusions of their papers agree or disagree with the professor's partisan preferences (*A* for Agree, *D* for Disagree), and whether or not their partisan ID matches their professor's (*M* for Match, *O* for Opposite). This creates four student types based on affinity and outcome characteristics (Table 1).

[Table 1: The Relationship Between Professor and Student Affinity and Outcome Characteristics]

13

Students labeled S_{MA} thus match their professor's party ID and write assignments that agree with their professor's partisan preferences, while students denoted S_{MD} match their professor in party ID but write papers that disagree with their professor's partisan beliefs. Students designated as S_{OA} have a partisan ID that is the opposite of their professor's, but who write papers that agree with the professor's partisan stance; students labeled S_{OD} identify with the party their professor opposes and write papers whose conclusions disagree with the professor's partisan beliefs.

We assume that a student's 'true' grade on a given topic, $g(S)$, is a function of student S's abilities and efforts, which are both general and assignment-specific. We leave the functional form of $g(S)$ unspecified. We do assume that the functional form is the same for all students, even though students' scores on each factor contributing to the outcome variable will obviously be different. Schematically, let us represent the formula as follows:

$$g(S) = \beta_1 Skill + \beta_2 Effort$$

We assume that the students have equal abilities and that they work equally hard, so that their 'true' (that is, unbiased) scores would be:

$$S_{MA} = S_{MD} = S_{OA} = S_{OD}.$$

Next, let us assume that the instructor's biases contaminate the grades, so that the grading function is now $\theta g(S)$ where θ denotes the professor's bias. For the unbiased professor, $\theta = 1$. If $\theta > 1$ the professor's biases favor the student, and if $\theta < 1$ the professor's bias harms the student. The more θ deviates from 1, the greater the bias.

Outcome Bias

14

We first consider 'outcome' bias, in which the professor favors *conclusions* that agree with the professor's partisan preferences and disfavors conclusions that disagree with them. We accordingly assume that $\theta_A > 1$ and $\theta_D < 1$.

Under these assumptions, if outcome bias was the only form of grading bias, then $\theta_A S_{MA} = \theta_A S_{OA} > \theta_D S_{MD} = \theta_D S_{OD}$ and all essays would receive biased scores, with those agreeing with the professor's preferred outcomes receiving higher scores than those disagreeing with the professor's outcome preference.

If the professor required that the students write in favor of a specific outcome (whether *A* or *D*), then all students would receive biased grades, but the rank ordering of students would not be affected. If *all* students wrote essays agreeing (or disagreeing) with the professor's outcome preference for legalizing same sex marriage, for example, all grades would be biased upward (or downward) but no student would be disadvantaged relative to the student's peers.[8] But if the professor allowed the students to choose the outcome that they write about, and if some students write essays in favor of legalization and some against it, then the differential impact of the bias on the students is obvious. Those who write papers that agree with the professor's outcome preferences would get higher scores than they deserve; the others, lower scores.

It is worth noting that if students believe that their professor manifests outcome bias, their incentives will be distorted. Students who agree with the professor would be discouraged from taking the hard task of writing a paper that challenges those views, while students who disagree with the professor face incentives to adopt the professor's perspectives in order to improve their grades. Outcome bias is likely to induce student papers that pander.

It is not possible to determine whether outcome bias exists merely by comparing the scores of the students who write papers that agree with the professor's partisan preferences to the scores of

15

students who write papers that disagree with this perspective, as the differences might be attributable to the differences in skills and efforts between the two groups of students. A better test of this bias would be to compare the scores of students who write papers that agree with the professor's outcome preferences *as well as* papers that disagree with these preferences. If all students wrote both types of papers, then the difference between the mean scores of the two groups could not easily be attributed to differential skills or efforts.[9] Differences in the mean scores of the Agree and Disagree papers would more likely reflect differential treatment by the professor in grading: that is, political outcome bias.

Affinity Bias

The second form of political grading bias is affinity bias, where the professor illegitimately favors *students* who share the professor's partisan allegiances, independent of the conclusions they reach. We identify the bias as $\theta_M > 1$ and $\theta_O < 1$ where the subscript M denotes a partisan match and O a partisan mismatch, and which implies that $\theta_M S_{MA}, \theta_M S_{MD} > 1 > \theta_O S_{OA}, \theta_O S_{OD}$. If affinity bias were the *only* form of political bias, then $\theta_M S_{MA} = \theta_M S_{MD} > 1 > \theta_O S_{OA} = \theta_O S_{OD}$ and the difference in grades between the two students would depend entirely on whether or not the student's partisanship matched the professor's or not.

Affinity bias might stem from any one of several sources. In its most unfair version, the instructor might consciously decide to punish partisans of the other side and reward co-partisans by cataloging students' beliefs and identifications. We believe that it would be professionally irresponsible and unethical for a professor to collect data on students' partisanship and to use these data in assigning grades, and we hope this does not happen. A similar process might happen unconsciously, as a professor could inadvertently develop knowledge of students' partisanship through course interactions, and use this information in assigning (biased) grades:

16

yet another reason to have written work submitted anonymously. It is plausible that a professor could infer the partisanship of the students through textual clues from the assignments themselves, a possibility that should be taken seriously. Finally, if the perceptual-filters mechanism is at work—that is, if a liberal professor and a liberal student both have a shared, but skewed, view of what a "good" and "bad" conservative essay should be, a classification not shared by conservatives—then affinity bias can creep in without any overt or inadvertent classification of students.

Determining whether affinity bias exists is tricky, as it requires comparing the mean scores of those students who match the professor's party ID with those who do not. Even if such data were available, they would be insufficient if students with different party IDs wrote papers that arrived at differing conclusions. In this case, if the scores of S_{MA} were higher than S_{OD}, the gap could be due to affinity bias or outcome bias. In order to learn whether affinity *differences* exist, then it is necessary for the professor to know (or at least be able to guess) the partisanship of the students *and* to have the students write papers that agree and disagree with the professor's preferences. Only by comparing S_{MA} and S_{MD} to S_{OA} and S_{OD} can we see whether the mean differences in scores are linked to partisan affiliations.[10]

Without a dataset linking students' scores to student partisanship, it would be impossible to learn whether affinity differences exist. Determining whether affinity bias exists requires experimental data that compares students' party IDs and paper conclusions to those of their professors. In a separate paper we have obtained and analyzed such data (Author 2012), so we focus our empirical analysis below exclusively on outcome bias. Note, however, that even in the presence of affinity bias, the logic of paired assignments described below will still result in equal or less bias than in non-paired assignments.

17

THE LOGIC OF PAIRED ASSIGNMENTS

One way to reduce the possibility of outcome bias is to 'pair' assignments so that students write one paper that concludes 'Agree' while the other paper concludes 'Disagree', so that the student's score for the assignment is the average of the two papers.

To see this, remember that outcome bias implies that $\theta_A > 1$ and $\theta_D < 1$ and that, under pure outcome bias, $\theta_A S_{MA} = \theta_A S_{OA} > \theta_D S_{MD} = \theta_D S_{OD}$ so that all individual papers receive biased scores. (If there were no bias at all, then $\theta_A = \theta_D = 1$.) If the professor requires the paired assignments in which the students write both the Agree and Disagree essays, however, the overall grade for each student will be $(\theta_A S_{MA} + \theta_D S_{MD})/2$ for the student who shares the professor's partisanship and $(\theta_A S_{OA} + \theta_D S_{OD})/2$ for the student whose partisan ID is different from the professor. If the professor's outcome bias is symmetrical, so that abs$[\theta_A = \theta_D]$, then in averaging the two scores the bias would be completely eliminated. Yet even if the bias is asymmetrical, with the professor showing greater bias toward preferred outcomes than against them (or vice versa), the paired essays would bring each student's overall grade closer to the unbiased grade and, as importantly, the pairing would create a proper ordinal ranking of the students.

Consider, for example, a situation where $\theta_A = 1.1$ (that is, a 10 percent bias in favor of papers that agree with the professor's preferences and $\theta_D = 0.8$ (so that the student disagreeing with the professor has a 20 percent bias penalty). Assume two students are identical in abilities and efforts, so that their unbiased grades would be equal. If one student wrote a paper that agreed with the professor and the other student a paper that disagreed, the student scores would be $(1.1)S_A > (0.8)S_D$. With matching, the scores for both would be $(1.1 + 0.8)/2 = 0.95S$. The overall scores would still be biased, but no students would be favored or disfavored.

18

On the other hand, if affinity bias were the *only* form of political bias (so that $\theta_M S_{MA} = \theta_M S_{MD} > 1 > \theta_O S_{OA} = \theta_O S_{OD}$) then pairing assignments would not reduce bias as $(\theta_M S_{MA} + \theta_M S_{MD})/2 = \theta_M$ and so forth. Yet if affinity bias exists in conjunction with outcome bias (so that $\theta_M S_{MA} > \theta_M S_{MD} > 1 > \theta_O S_{OA} > \theta_O S_{OD}$), then pairing would still reduce the overall bias in scoring, as for example $(\theta_M S_{MA} + \theta_M S_{MD})/2$ would be closer to 1 than $\theta_M S_{MA}$ alone.

Creating appropriately paired assignments reduces the possibility of outcome bias, at least so far as it affects students' overall grades. Moreover, it allows professors to check *post hoc* whether their scores did reveal outcome bias. This information can be used by professors for self-reflection, or it can be revealed to the students to enhance the transparency of the professors' grading schemes (with the further disclosure that the pairing eliminated the outcome bias). The paired assignment strategy can work even when, by all accounts, one recommendation is clearly 'superior' to the other. Consider an extreme case: "Should the United States reinstitute slavery?" It would be difficult for any modern American to write a defense of slavery that is as compelling as an attack on the institution, but let us assume that a few students actually are convinced that slavery is the preferred option. By pairing the assignments, these students would not face outcome bias, because each student would be assessed on the quality of both pro and con arguments. A further benefit of the paired assignments is that they reward the students who think most carefully, and write most clearly, on the differing sides of the issue.

Pairing assignments can be useful on virtually any political topic, and it need not involve a normative matter. For example, a professor could ask questions such as "Are voters rational?", "Does 'going public' enhance presidential power?", "Are the media politically biased?", and so forth. Even if the professor has strong, empirically-based convictions on these questions, the

19

pairing helps ensure that students who agree or disagree with the professor are neither advantaged nor disadvantaged.

AN EMPIRICAL EXAMINATION OF OUTCOME BIAS

Comparing grades of students within and across classes can do little to determine whether political outcome bias exists. A better test would compare the grades of the *same* students on assignments in a required class for which they must write Agree and Disagree papers on the same topic. If the course is required, then selection bias within majors is eliminated.[11] If students have incentives to obtain high grades, then the differences between the average scores of the Agree and Disagree papers across the students should be an unbiased indicator of the professor's outcome bias. Differences in the preparation, motivation or attitudes of the students would be reflected in each student's total score, but these student characteristics should not discriminate in favor of either Agree or Disagree papers. Student perception of their professors' ideology should be inconsequential, as the students would write papers on which they can assume that the professor himself has both Agree and Disagree positions.

Fortunately, [Author] has such data from the "Ethics and Values in Public Policy" course he routinely teaches. The class typically enrolls between 12 and 20 graduate students, it is a required course in the [anonymous program], and it inevitably covers controversial political and moral issues. Over the past several years [Author] has tinkered, in an *ad hoc* way, with various assessment schemes, and has kept reasonably good records on topics and scores. The data do not come from true experiments (e.g., carefully designed with clear hypotheses, adequate controls, and so forth) and therefore cannot always answers the questions worth asking, but they can reveal some information on the potential (and actual) biases in grading.

20

Before considering these data, a couple points about investigating bias in one's own course must be revealed. [Author] hopes that he has not been biased, and believes that he tried to grade fairly. Still, he recognizes the possibility that he may have been biased, albeit inadvertently. He also acknowledges that, given the limitations of his data, he could have been biased in other semesters, or in other courses. As we are assessing the data from a single course with a single professor, whatever is revealed here cannot easily be generalized to other professors, who might be more or less biased than [Author] is. Finally, we recognize that critics can (and should) look skeptically at the results. If bias is not found, partisan critics of the academy might respond that of course professors who study themselves conclude they are unbiased. If bias is identified, partisan defenders of universities might infer that I found bias deliberately in order to support the conservative critique. Readers are likely to bring their own confirmation biases to this study.

We assess both partisan and 'preference' outcome bias. For partisan bias, our concern is whether papers whose conclusions agree with the party ID of the professor receive higher average scores than those that do not. For preference bias, we examine whether papers with conclusions that agree with the professor's policy preferences receive higher scores than those that disagree with these preferences, even if the preferences are not clearly linked with either the Democratic or Republican party.

The Data

The data come from scores [Author] gave to policy memos administered between 2007 and 2010. In each semester the students wrote two or three sets of policy memos, each based on a small list of possible topics. The students had to write two memos on each topic, with one memo recommending "Do X" and the other "Do Not Do X". The memos were submitted anonymously and [Author] recorded the position recommended and the score for each memo. A total of 156

paired memos were used in the analysis after unusable scores were eliminated.[12] Of these, 88 memos had clearly defined partisan elements that distinguished between Democratic and Republican positions; the remainder were not neatly divided along partisan lines.

Partisan Outcome Bias

To investigate whether partisan outcome bias exists, we calculated a paired-sample difference of means test comparing 'Republican' and 'Democratic' scores. The results from these tests appear in Table 2. A two-tailed test of significance would be appropriate if the partisanship of the professor was unknown, but as [Author] is a Democrat, confirmation bias suggests that the [Author] would be inclined to treat Democratic answers more sympathetically, suggesting a one-tailed test).

[Table 2: Difference of Means Test Comparing Republican and Democratic Scores about here]

The data indicate a 0.44-point difference in the mean scores received by Republican and Democratic papers, with the Republicans receiving the slightly lower scores on average. These differences were not significant in a two-tailed test ($t = 1.61$, $p = 0.11$), although they approached significance in a one-tailed test ($p = 0.056$).

One might conclude that [Author's] grading scheme does not reflect partisan outcome bias because the differences in scores are relatively modest in size and of borderline statistical significance But caution is in order. First, even fairly small differences can have a substantial impact on final grades (see Rom 2011). If student scores are distributed normally, a substantial number of students will be very close to the grading cut points; a half a point can easily make the difference between a B+ and an A-, for example.[13] We suspect that virtually all professors have

22

had students willing to appeal that B+ in seeking an A-, so what might seem a small difference to the professor can loom large to the student. Second, all else being equal, if the sample were bigger the differences would be statistically significant at the conventional levels.

Preference Bias

[Author's] political views are not always neatly aligned with a particular political party, and he acknowledges policy preferences on matters that do not distinguish Democrats from Republicans. To determine whether these policy preferences biased his grades, we coded whether papers agreed or disagreed with [Author's] preferences.[14] Table 3 presents the results of another paired sample t-test. Once again, these differences were not significant in a two-tailed test ($t = 1.60$, $p = 0.11$), although they approached significance in a one-tailed test ($p = 0.055$). As with the partisan bias test, the observed difference between [Author's] Agree and his Disagree scoring was less than 0.5 points.

[Table 2: Difference of Means Test Comparing Papers that Agree With Instructor's Preferences with Those That Disagree about here]

The differences in scores between groups based on partisan or preference outcomes were again small, but not necessarily inconsequential as noted above. Pairing the assignments would nevertheless negate any outcome bias that would have favored students coming to conclusions favored by the professor. Whatever outcome biases [Author] had were mitigated by the paired assignments.

OVERCOMING OUTCOME BIAS IN GRADING

Clear standards, randomizing papers, and anonymous submissions do not prevent political outcome bias.[15] Other measures are needed. One obvious solution to the bias problem is for

23

instructors to simply try to be fair. Such efforts are necessary, but not sufficient. Merely being conscientious "does not, as a rule, immunize one's judgments from the effects of bias" (Archer and McCarthy 1988, citing Nisbett and Ross 1980). Training is no panacea either: if the biases are deeply entrenched, they can be difficult to overcome them (Sweedler-Brown 1992).

Grading and the Ethic of Responsibility

Preventing outcome bias requires a mixture of both an 'ethic of responsibility' and an 'ethic of rules' (see Weber 1946: 77-129). For the former, the professor needs to recognize the possibility of outcome bias and to consider thoughtfully the ways to avoid it. For the latter, the professor need establish rules designed to mitigate political bias. As professors – understandably – are likely to resist having rules imposed on them concerning political bias (see, for example, Larkin 2004) and as 'good faith efforts' alone are unlikely to be effective, the ethic of responsibility should lead them to develop their own rules (in the course syllabus, for example) indicating what steps they are taking to avoid bias.[16]

The ethic of responsibility suggests that a professor should avoid giving assignments for which the professor knows that it is difficult, if not impossible, to remain unbiased regarding the outcomes. This can be harder than it looks, given that professors can have the strongest convictions on topics that are central to their courses and on which they have substantial expertise. Given the likelihood of confirmation bias, these convictions make it unlikely that professors can neutrally grade arguments that challenge or contradict the professors' biases.

If professors hold strong convictions on issues central to the courses they teach, how can assignments be designed to avoid outcome bias in grading? One important condition is that the professor must look into the mirror: "Do I have any reason to question my own impartiality, to suspect that I might be prejudiced for or against the [case] for any reason?" (Quote from judge

Thomas Penfield Jackson cited by Ayers 1990). This is a difficult task, but we should expect at

least as much introspection and honesty from a professor in the classroom as a citizen in the courtroom. If the answer is "I cannot be fair" then the potential assignment *must* be modified. One possibility, if a uniform assignment is given to the entire class, is to modify the assigned topic so that the instructor does *not* have strong convictions regarding the answers. For instance, [Author] has given the assignment "Should policies regarding same sex marriage be determined by the courts, or by voters?" Although [Author] has firm scholarly and personal commitments regarding the topic of same sex marriage – he believes that the case for same sex marriage is convincingly superior to the case against it -- he remains ambivalent on the answers to the procedural question. Judicial rulings are undemocratic, arguably ineffectual (see Rosenberg 1991), and potentially counterproductive (Diller 2000; Carr and Krieger 2003), yet perhaps necessary to protect minority rights (Yarbrough n.d.). And although democratic consent ultimately may be necessary for society to come to peace on an issue, voters can be intolerant of minorities. By creating assignments for which the professor truly has few preconceptions, the possibilities of unbiased assessments should increase.

Grading and the Ethic of Rules

As important as an ethic of responsibility is, such a moral turn may not be sufficient to protect against outcome bias. Thus it is imperative that the professor design a grading scheme that safeguards neutrality by design. Since normative questions ultimately lie at the heart of many classes, for undergraduate and graduate students alike, it is nearly impossible and clearly undesirable to simply drop all assignments that are not fully 'objective.'

Designing an unbiased grading scheme involves two main components. First, the syllabus should make clear that grading fairly is the instructor's priority, and the syllabus should outline what

steps will be taken to ensure fairness. The syllabus should note, for example, that written work will be submitted by student ID, how assignments should be formatted, what (if any) penalties are incurred for late assignments, and how concerns about bias should be addressed. Moreover, the assignments themselves should be designed to minimize bias, and these design elements should be made transparent to the students. Incorporating paired assignments, as well as the justification for them, into the syllabus has the potential to reduce political outcome bias in practice as well as perception.

CONCLUSIONS

Political bias in grading is both a hot topic and a hot potato. Political conservatives believe that political bias in academia is a serious problem. The authors believe that the problem of political bias should be taken seriously and that the academy should treat it with appropriate gravity. It is essential for faculty to be aware that bias is possible and to take steps to reduce the possibility that it will occur in fact or perception. Doing so benefits the students and the credibility of the faculty. Fortunately, in many situations, a simple tweak to assignments can reduce the possibility that bias will affect students' grades while also providing obvious transparency to outside critics. Pairing assignments will have the potential to mitigate if not eliminate partisan outcome bias. We highly recommend their use.

ENDNOTES

[1] Thanks to Ashley Branca and Rose Tutera-Baldauf for providing research assistance on this project and to Patrick Carr for his helpful comments. Thanks to Rachel Blum for proofreading the paper. Thanks also to the participants in the various seminars in which preliminary ideas were presented and also to two anonymous reviewers for *JPSE*.

[2] Halo and horn effects are a variant of confirmation bias, as discussed below.

[3] It is always possible for the grader to identify the student before grading the assignment, so the effectiveness of this approach depends on the instructor grading first and identifying later. In our view, doing so is consistent with professional norms of grading the assignment, not the person.

[4] Although partisan and ideological preferences are not necessarily the same, for the rest of this paper we will use the shorthand 'partisan' to mean both 'partisan and ideological'.

[5] We appreciate that partisan identifiers do not necessarily agree with all their party's policy positions, so these examples are illustrative rather than definitive. We assume that, in this case, Republicans sincerely believe that voter fraud is the key issue and that Democrats sincerely believe that voter suppression is the central concern.

[6] An anonymous reviewer of an earlier version of this paper argued that a professor could be biased towards outcomes as well as types of arguments. Using the notation developed below, we define the students as S_{AS}, S_{AN}, S_{DS}, and S_{DN} where the A and D subscripts indicate whether the student agrees or disagrees with the professor's partisan position. The S subscript denotes 'strong' essays that support that student's partisan preferences and N denotes 'nuanced' essays that disagree with that student's partisan position.

This creates the bias $\theta_{AS} > 1 > \theta_{DS}$ and possibly – although the reviewer does not claim this -- $\theta_{AN} > 1 > \theta_{DN}$. If the professor assigns a single essay in which the student is allowed to choose the conclusion, and students write sincere essays in which they defend their true beliefs, then the bias is direct: students who agree with the professor receive scores that are biased upward, and students who disagree with the professor receive scores biased downward. This is actually another form of outcome bias: the student who agrees with the professor receives a higher score than deserved, and the student disagreeing with the professor receives a lower score than warranted. Pairing assignments would reduce the bias under such a scenario.

[7] Google scholar reports that this article has been cited 25 times, but none of the subsequent articles directly examine political bias in grading. A discussion of this article in the *Chronicle of Higher Education* (Jacobson 2006) received extension reader commentary describing personal anecdotes involving perceived bias or fair treatment. Some studies have also found that liberals and conservatives have "differences in general neurocognitive functioning" (Amodio *et al.* 2007, 1; see also Jost *et al.* 2003). Others have disputed whether these results are relevant (Saletan 2007). Another study provides evidence of the complicated relationship between political orientation and cognitive ability (Kemmelmeier 2008).

[8] It would seem quite odd for a professor to specify that student papers must all produce the same conclusions.

[9] It is possible that students would generally devote more effort to writing papers that agree (or disagree) with their own preferences, although it seems reasonable that if students are motivated to earn high grades they would expend roughly equal effort to both tasks.

[10] We note that even if these differences exist it would not be proof of bias, as it is possible that students of one party ID actually are better than students of the opposing party ID. But claiming that Democrats are better students than Republicans, or vice versa, would be a highly controversial proposition.

[11] Selection bias would still affect the students' choice of their major, but it would not affect those within a given major.

[12] Scores were eliminated if [Author] did not record the positions or link scores to a particular position. A list of the memo topics is available from the Authors upon request.

[13] Fortunately, because each student wrote both a 'Republican' and 'Democratic' memo, whatever bias existed would not serve to advantage either Republican or Democratic students.

[14] In assigning [Author's] preferences, we followed the following decision rule: he uniformly voted with the Democrats on partisan issues, he identified his clear preference on other issues, and we forced him to vote on the issues on which he seemed to be ambivalent.

[15] Although, as noted, anonymous submissions can reduce the possibility of affinity bias.

[16] Students can generally appeal their grades if they believe the professor is biased, but the norms of faculty autonomy make such challenges quite difficult.

TABLES

Table 1: The Relationship Between Professor and Student Affinity and Outcome Characteristics

	Affinity Characteristics	
Outcome Characteristics	Match	Opposite
Agree	S_{MA}	S_{OA}
Disagree	S_{MD}	S_{OD}

Table 2: Difference of Means Test Comparing Republican and Democratic Scores

Variable	*N*	*Mean*	*Std. Dev.*
Democratic Essays	88	92.81	2.78
Republican Essays	88	92.38	2.86
Difference	88	0.44	2.56
H_0: Difference of Means ≥ 0		t = 1.61	p = 0.056
H_0: Difference of Means = 0			p = 0.11

Mark Carl Rom 9/17/12 1:40 PM
Deleted: 36

Mark Carl Rom 9/17/12 1:40 PM
Deleted: =

Table 3: Difference of Means Test Comparing Papers that Agree With Instructor's Preferences with Those That Disagree

Variable	N	Mean	Std. Dev.
Agree with Instructor	156	92.75	2.97
Disagree with Instructor	156	92.41	3.27
Difference	156	0.34	2.65
H_0: Difference of Means $\geq$ 0		t = 1.60	p = 0.055
H_0: Difference of Means = 0			p = 0.11

30

REFERENCES

Abramowitz, Stephen I., Beverly Gomes and Christine V. Abramowitz. 1975. "Publish or
Politic: Referee Bias in Manuscript Review." *Journal of Applied Social Psychology*. 5(3):
187-200.

ACTA (American Council of Trustees and Alumni). 2006. *How Many Ward Churchills?*
Website available at https://www.goacta.org/publications/downloads/ChurchillFinal.pdf.

Amodio, David M., John T. Jost, Sarah L. Master and Cindy M. Yee. 2007. "Neurocognitive
Correlates of Liberalism and Conservativism." *Nature Neuroscience*. September. Website
available at
http://www.psych.nyu.edu/amodiolab/Amodio%20et%20al.%20(2007)%20Nature%20N
euro.pdf.

Anonymous. 2006. Posting on RateMyProfessors.com. Website available at
http://www.ratemyprofessors.com/ShowRatings.jsp?tid=134563&page=2.

Author. 2012. "Political Affinity Bias in Grading." Unpublished manuscript.

Archer, J., and B. McCarthy. 1988. "Personal Biases in Student Assessment." *Educational
Research* 30:142–45.

Ayers, B. Drummond, Jr. 1990. "Judge Urges Impartiality by Jurors As Capital Mayor's Trial
Begins." *New York Times*. June 5. Website available at
http://query.nytimes.com/gst/fullpage.html?res=9C0CE6DF1130F936A35755C0A96695
8260.

Balch, Stephen. 2006. "The Study Needs Study." Blog posting. *National Review Online: Phi Beta Cons*. March 30. Website available at http://phibetacons.nationalreview.com/post/?q=ODJjM2U1YTg5M2VkNjYzZGI4NGIzODFiYWJiZGFhOTk=.

Bar, Talia and Asaf Zussman. 2012. "Partisan Grading." *American Economic Journal: Applied Economics*. 4(1): 30-48.

Belfield, Clive and Henry M. Levin. 2005. "Vouchers and Public Policy: When Ideology Trumps Evidence." *American Journal of Education*. August. 111: 548-567. Website available at http://www.journals.uchicago.edu/doi/abs/10.1086/431183.

Carr, Karen Eva and Linda Hamilton Krieger. 2003. *Backlash Against the ADA*. Ann Arbor: University of Michigan Press.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. New York: John Wiley & Sons, Inc.

Ceci, Stephen J., Douglas Peters, and Jonathan Plotkin. 1985. "Human Subjects Review, Personal Values, and the Regulation of Social Science Research." *American Psychologist*. 40(9): 994-1002.

Cohen-Cole, Ethan and Steven Durlauf. 2005. "Evaluating Claims of Bias in Academia: A Comment on Klein and Western's 'How Many Democrats per Republican at UC-Berkeley and Stanford'." Working paper.

Daly, J.A. and F. Dickson-Markman. 1982. "Contrast Effects in Evaluating Essays." *Journal of Educational Measurement*. 19(4): 309-16.

Diller, Matthew. 2000. "Judicial Backlash, the ADA, and the Civil Rights Model." *Berkeley Journal of Employment and Labor Law.* 21(1): 19-53.

Fleming, Neil D. 1999. "Biases in Marking Students' Written Work: Quality?" *Assessment Matters in Higher Education: Choosing and Using Diverse Approaches.* Sally Brown and Angela Glasner, editors. Buckingham: Society for Research into Higher Education and Open University Press.

Gross, Neil and Solon Simmons. 2007. "The Social and Political Views of American Professors." Working paper. September 24. Website available at http://www.wjh.harvard.edu/~ngross/lounsbery_9-25.pdf.

Haidt, Jonathan. "Post-Partisan Social Psychology." Website available at http://people.virginia.edu/~jdh6n/postpartisan.html.

Hales, L.W. and E. Tokar. 1975. "The Effect of the Quality of Preceding Responses on the Grades Assigned to Subsequent Responses to an Essay Question." *Journal of Educational Measurement.* 12(2): 115-17.

Horowitz, David. 2007. *Indoctrination U: The Left's War Against Academic Freedom.* New York: Encounter Books.

Horowitz, David and Eli Lehrer. 2002. "Political Bias in the Administrations and Faculties of 32 Elite Colleges and Universities." Center for the Study of Popular Culture, Los Angeles.

Hughes, D.C., B. Keeling, and B.F. Tuck. 1980. "The Influence of Context Position and Scoring Method on Essay Scoring." *Journal of Educational Measurement.* 17(2): 131-35.

33

Gary Smith 6/11/13 12:42 PM
Comment [4]: Is there a date available for this source?

Inbar, Yoel and Joris Lammers. 2012. "Political Diversity in Social and Personality Psychology."
Social Science Research Network. Website available at
http://ssrn.com/abstract=2002636.

Jacobson, Jennifer. 2006. "Study Casts Doubt on Claims that Conservative Students Face
Discrimination in Classes" *Chronicle of Higher Education.* March 30. Reposted, with
comments at http://www.freerepublic.com/focus/f-news/1606178/posts.

Jaschik, Scott. 2006. "The Real Bias in the Classroom." *Inside Higher Ed.* March 20. Website
available at http://www.insidehighered.com/news/2006/03/20/politics.

Johnson, L. 1996. "Resisting Change: Information-Seeking and Stereotype Change." *European
Journal of Social Psychology*. 26: 799–825.

Jonas, Eva, Stefan Schulz-Hardt, Dieter Frey, and Norman Thelen. 2001. "Confirmation Bias in
Sequential Information Search after Preliminary Decisions: An Expansion of Dissonance
Theoretical Research on Selective Exposure to Information." *Journal of Personality and
Social Psychology*. 80: 557–571.

Jost, John T., Jack Glaser, Arie W. Kruglanski, Frank J. Sulloway. 2003. "Political
Conservativism as Motivated Social Cognition." *Psychological Bulletin*. 129(3): 339–375
Website available at http://faculty.virginia.edu/haidtlab/jost.glaser.political-conservatism-
as-motivated-social-cog.pdf.

Kelly-Woessner, April and Matthew C. Woessner. 2006. "My Professor is a Partisan Hack: How
Perceptions of a Professor's Political Views Affect Student Course Evaluations." *PS:
Political Science & Politics*. (July) 495-501.

Kemmelmeier, Markus, Cherry Danielson, and Jay Basten. 2005. "What's in a Grade? Academic
Success and Political Orientation." *Personality and Social Psychology Bulletin*. 31(10):
1386-1399. Website available at http://psp.sagepub.com/content/31/10/1386.short.

Kemmelmeir, Markus. 2008. "Is There a Relationship Between Political Orientation and
Cognitive Ability? A Test of Three Hypotheses in Two Studies." *Personality and
Individual Differences*. 45(8): 767-72. Website available at
http://www.sciencedirect.com/science/article/pii/S0191886908002912.

Kennedy, Peter. 2001. *A Guide to Econometrics*. Cambridge, MA: MIT Press.

Klein, Daniel and Charlotta Stern. 2004-5. "Political Diversity in Six Disciplines." *Academic
Questions* 18:40-52.

Klein, Daniel and Andrew Western. 2004-5. "Voter Registration of Berkeley and Stanford
Faculty." *Academic Questions* 18:53-65.

Larkin, Graham. 2004. "What's Not to Like about the Academic Bill of Rights?" Website
available at http://www.aaup-ca.org/Larkin_abor.html.

Lee, John. 2006. "The 'Faculty Bias' Studies: Science or Propaganda?" American Federation of
Teachers, Washington, D.C. Website available at
http://www.freeexchangeoncampus.org/index.php?option=com_docman&task=cat_view
&gid=22&Itemid=25.

Leef, George. 2006. "Re: Classroom Bias." *National Review Online: Phi Beta Cons*. March 30.
Website available at
http://phibetacons.nationalreview.com/post/?q=YzI5ZjJmMmNmM2YwZTg1NGNmOT
g2ZDJmOGFhNjI3ODc=.

35

Lord, C. G., L. Ross, and M. R. Leeper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology* 37: 2098–2109.

Lundgren, S. R., and R. Prislin. 1998. "Motivated Cognitive Processing and Attitude Change." *Personality and Social Psychology Bulletin*. 24: 715–26.

McCardle, Megan. 2011a. "Unbiasing Academia." *TheAtlantic.com* February 8. Website available at http://www.theatlantic.com/national/archive/2011/02/unbiasing-academia/70955/.

McCardle, Megan 2011b. "What Does Bias Look Like?" *TheAtlantic.com*. February 14. Website available at http://www.theatlantic.com/national/archive/2011/02/what-does-bias-look-like/71153/.

Mahoney, Michael J. 1977. "Publication Prejudices: An Experimental Study of Confirmatory Bias in the Peer Review System." *Cognitive Therapy and Research*. 1(2): 161-175.

Mariani, Mack D. and Gordon J. Hewitt. 2008. "Indoctrination U.? Faculty Ideology and Changes in Student Political Orientation." *PS: Political Science & Politics*. 41(4): 773-783.

Munro, Geoffrey D., Terell P. Lasane, and Scott P. Leary. 2010. "Political Partisan Prejudice: Selective Distortion and Weighting of Evaluative Categories in College Admissions Applications." *Journal of Applied Social Psychology*. 40(9): 2434-62.

Nisbett, R., and L. Ross. 1980. Human Inference. Upper Saddle River, NJ: Prentice Hall.

Rom, Mark Carl. 2011. "Grading Accurately." *Journal of Political Science Education*. 7: 208-223.

Rosenberg, Gerald N. *The Hollow Hope: Can Courts Bring About Social Change?* Chicago: The

    University of Chicago Press.

Rothman, Stanley, S. Robert Lichter, and Neil Nevitte. 2005. "Politics and Professional

    Advancement Among College Faculty." Article 2 in *The Forum*, vol. 3.

Rothman, Stanley, April Kelly-Woessner and Matthew Woessner. 2011. *The Still Divided*

    Academy: How Competing Visions of Power, Politics, and Diversity Complicate the

    Mission of Higher Education. Lanham, MD: Rowman & Littlefield Publishers.

Saletan, William. 2007. "Liberal Interpretation: Rigging a Study to Make Conservatives Look

    Stupid." *Slate*. (September 14). Website available at http://www.slate.com/id/2173965/.

Smallwood, Scott. 2006. "Poll Finds the Public Opposes Government Interference in Academe."

    The Chronicle of Higher Education. June 16. Website available at

    http://chronicle.com/free/v52/i41/41a00101.htm.

Spear, M-G. 1984. The Biasing Influence of Pupil Sex in a Science Marking Exercise." *Research*

    in Science and Technological Education. 2(1): 55-60.

Students for Academic Freedom. N.d. "Academic Bill of Rights." Website available at

    http://www.studentsforacademicfreedom.org/documents/1925/abor.html.

Sweedler-Brown, C. 1992. "The Effect of Training on the Appearance Bias of Holistic Essay

    Graders." *Journal of Research and Development in Education*. 26(1): 24–29.

Tierney, John. 2011. "Social Scientist Sees Bias Within." *The New York Times*. February 7.

    Website available at

    http://www.nytimes.com/2011/02/08/science/08tier.html?_r=3&ref=science.

Gary Smith 6/11/13 12:47 PM
Comment [5]: There's a date provided in the in-text citation, but not here? Is this an N.D. or just an oversight?

Vitiello, Michael. 2007. "Liberal Bias in the Legal Academy: Overstated and Undervalued."

    Unpublished manuscript. Website available at

    http://works.bepress.com/michael_vitiello/1/.

Weber, Max. 1946. "Politics as a Vocation. " In *Max Weber*: *Essays in Sociology*. Gerth and

    Mills, editors. London: Oxford University Press.

Woessner, Matthew and Kelly-Woessner. 2009. "I Think My Professor is a Democrat:

    Considering Whether Students Recognize and React to Faculty Politics." *PS: Political*

    Science and Politics. 42(2): 343-52.

Yarbrough, Tinsley. N.d. "Protecting Minority Rights." U.S. Department of State, International

    Information Programs. Website available at

    http://usinfo.state.gov/products/pubs/democracy/dmpaper11.htm.

Zipp, John F. and Rudy Fenwick. 2006. "Is the Academy a Liberal Hegemony? The Political

    Orientations and Educational Values of Professors." *Public Opinion Quarterly* 70:304-

    326.