

ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

LAWYERS' ROLE-INDUCED BIAS ARISES FAST AND PERSISTS DESPITE INTERVENTION

Holger Spamann

Forthcoming in *Journal of Legal Studies*

Discussion Paper No. 1005
Revision

06/2020

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center/

The Social Science Research Network Electronic Paper Collection:
<https://ssrn.com/abstract=3387536>

Lawyers' Role-Induced Bias Arises Fast and Persists Despite Intervention

Holger Spamann *

Abstract

Law students randomly assigned to represent one side in a legal argument in the classroom exhibit substantial role-induced prediction bias for their side within only 40 minutes of their role assignment. Reminding students that prediction requires a more neutral perspective than advocacy does not attenuate the bias. The bias occurs evenly in male and female participants, who also report equal confidence in their predictions.

* hspamann@law.harvard.edu. This research uses data from a classroom exercise originally collected for pedagogical purposes, and on these grounds was determined not to be human subject research by Harvard IRB19-0611. For very helpful comments, I thank two anonymous referees and the editor, William Hubbard. The data and code generating this paper's statistics and figures are available at <https://doi.org/10.7910/DVN/CRZCPT>.

1 Introduction

Lawyers need to make predictions about the likely success of legal arguments to advise their clients and to plan steps in litigation and transactions. It is well documented in general psychology, however, that humans tend to overestimate the strengths of their own side's arguments, presenting a threat to lawyers' effectiveness at least in their role as advisors. Can law school and other professional training overcome this role-induced bias by fostering awareness of competing perspectives? Eigen & Listokin (2012) show that the answer is no at the end of a moot court competition, where law students' median time investment was over 30 hours in the course of several weeks. But lawyers' predictions may be more important at the beginning of a case. What if lawyers were asked in the first hour of their engagement with the case? Loewenstein et al. (1993) and Babcock et al. (1995) found that such short time was sufficient to induce bias, but their experiment revolved around a factual question and put law students in the shoes of litigants, not advocates. What if time were short, the question were legal, *and* law students acted as advocates? And if bias emerged in such conditions, could a simple intervention eliminate it?

This short paper shows that role-induced bias does arise fast in advocates and survives a simple intervention, at least in an experiment with law students. After only 40 minutes or less of engagement with a case, law students exhibit substantial role-induced bias towards the side they have been randomly assigned to represent in the classroom. Of students randomly assigned to represent petitioner or respondent in a mock Supreme Court oral argument of *U.S. v. Newman*¹, over 50% of those assigned petitioner's counsel's role predicted petitioner victory, as opposed to less than 20% of those assigned respondent's counsel's role, with randomly assigned justices in the middle. These differences in prediction were actually starker when the role assignment was performed a mere 30 minutes before the 10-minute mock argument than when the assignment was performed five days prior, before the students ever started reading about the relevant law. Reminding students before making their prediction that their role in the oral argument had been "that of an advocate, whereas we are now stepping back out and [...] asking your opinion as a neutral, sober observer" has no detectable effect on their biased predictions. Similar to the findings in Eigen and Listokin (2012), the bias occurs almost evenly in male and female participants, who also report equal confidence in their predictions.

The rest of this paper proceeds as follows. Section 2 briefly reviews the prior literature. Section 3 states the hypotheses. Section 4 presents the experimental design. Section 5 describes the data collected. Section 6 reports the results. Section 7 discusses limitations and implications and concludes. The experimental materials are provided as an online appendix, and all the data and code are available at <https://doi.org/10.7910/DVN/CRZCPT>.

2 Prior Literature

Eigen & Listokin (2012) and Engel & Glöckner (2013) provide excellent recent reviews of role-induced bias with special reference to law. The neutral term "role-induced bias" seems a better descriptor than "optimism bias," "self-serving bias," or "motivated reasoning" because as Engel & Glöckner (2013) point out, the bias persists even after individuals shed their role and hence any motive to distort their view, which is more consistent with coherence-based reasoning (e.g., Simon et al. 2015) or other variants of pre-decisional information distortion (cf. DeKay 2015; Russo 2018). From a practical perspective, little

¹ *U.S. v. Newman*, 773 F.3d 438 (2014), cert. denied, 136 S.Ct. 242.

hinges on the mechanism and the degree of persistence: either way, the deviation from Bayesian inference will lead to erroneous prediction and strategizing while a lawyer is engaged in a case.

Observational evidence suggests that role-induced bias is practically relevant for lawyers. Goodman-Delahunty et al. (2010) show that actual lawyers' predictions of outcomes in cases they work on are biased in the direction of their side. As Eigen & Listokin (2012) note, however, this observational evidence is confounded by lawyers' choices which side to work for in the first place.

The experimental study of role-induced bias in law starts with Loewenstein et al. (1993) and Babcock et al. (1995) (replicated in Hippel & Hoepfner 2019). In these experiments, students (including law students) randomly assigned the role of one party in a fact-centered dispute tended to overestimate the strength of their case after half an hour of preparation, leading to bargaining impasse. Babcock et al. (1997) found that this bias and the ensuing bargaining impasse could be eliminated by asking participants to write down the weaknesses of their case before negotiation. As Eigen & Listokin (2012, 244) point out, however, the roles of party and advocate are distinct and might well trigger different psychological processes. For similar reasons, the findings of political bias in the perception of legal arguments by law clerks, law students, and undergraduate students in Braman & Nelson (2007) and Furgeson et al. (2008a, 2008b) do not necessarily imply that lawyers would be biased in their perception of argument in disputes in which they do not have a stake and prior view, unlike in salient political questions.

Closer to the role of advocate is a design pioneered by Simon et al. (2008) (building on Simon et al. 2004) and recently extended by Engel & Glöckner (2013). Simon et al. (2008) put subjects in the role of an investigator charged with investigating the facts in a university disciplinary procedure on behalf of one the two sides. Engel & Glöckner (2013) asked subjects to sketch the pleadings in a criminal case as an intern for the prosecutor or defense counsel. The subjects were lay people and the questions factual, however, whereas U.S. law schools place heavy emphasis on teaching the indeterminacy of much legal argument, which could better calibrate law students' and lawyers' predictions, in particular with respect to *legal* questions.

Lawyers as advocates in legal disputes are studied in Eigen & Listokin (2012), who survey law students at the end of a moot court at four U.S. law schools. At three of the four schools, role "assignments were effectively random" "[t]o the best of [the authors'] knowledge" (248). As already noted, however, the law students' median time investment in the case before completing the survey was over 30 hours spaced out over several weeks. This leaves open the main question studied in this paper: Does short engagement with a case (as in Loewenstein et al. 1993; Babcock et al. 1995; and others) as a neutral advocate in a legal dispute (as in Eigen & Listokin 2012) create role-induced bias in law students, and if so, can this bias be turned off with a simple debiasing intervention?

An ancillary question studied in this paper is whether role-induced bias is more pronounced in males than in females. Goodman-Delahunty et al. (2010) find that it is, while Eigen & Listokin (2012) do not. The answer is of political interest because it might help assess the social costs and causes of female underrepresentation in the legal profession (cf., e.g., European Parliament 2017). Theoretically, one might expect males to be more biased if, like Eigen & Listokin (2012), one conceives of role-induced bias as a form of overconfidence. "Perhaps the most high-profile claim" in the literature on individual differences in overconfidence "has been that men show more overplacement than women" (Moore & Dev 2018). That said, overplacement—"the exaggerated belief that one is better than others"—is a different form of

overconfidence than the one possibly implicated in role-induced bias, and even “male overplacement is far from universal ..., and plenty of studies have failed to find gender differences” (Moore & Dev 2018).

3 Hypotheses

As a preliminary matter, note that bias is not defined with respect to a single prediction when the truth is unknown and possibly unknowable. However, since the true probability distribution is necessarily unique, differences in predictions do identify bias because at least one prediction must be wrong. Role-induced bias is thus identified by the systematic difference in predictions in the direction of the randomly assigned side.²

This leads to the following main hypotheses:

Main hypothesis 1 (Role-induced bias): Lawyers will be more likely to predict a win of the side they were randomly assigned to represent.

Main hypothesis 2 (Debiasing): Reminding lawyers that their role as advocate may interfere with their neutrality will reduce the role-induced differential in predictions.

These two main hypotheses will be tested experimentally using random assignment of the party to represent (main hypothesis 1) and cross-randomization of a reminder about the need for neutrality (main hypothesis 2). In addition, the data collected allow some inference about the following two ancillary hypotheses:

Ancillary hypothesis 1 (Mediating effect of time): The longer lawyers represent one side, the stronger they will be biased in its direction.

Ancillary hypothesis 2 (Mediating effect of gender): Females will be less biased than males.

Ancillary hypothesis 1 follows directly from the motivation of this paper as an extension of Eigen & Listokin (2012). However, the test here is not randomized (i.e., genuinely experimental) and hence subject to possible confounding by other factors, as explained below. With respect to ancillary hypothesis 2, the test is necessarily non-experimental because, like any immutable trait, gender cannot be randomly assigned. Motivation for ancillary hypothesis 2 is given at the end of section 2.

Naturally, no single paper can test these hypotheses in their generality. Outcomes might be dependent on the legal question at issue, the type of dispute, the subject pool, or the nature of the debiasing. The next section describes the setting and manner in which these hypotheses were tested here.

4 Design

4.1 Setting

The experimental data were generated as a byproduct of a classroom exercise. The class is the basic class on corporations/business organizations taught by the same instructor with identical materials at two top-

² Note that when defined in such relative terms, the bias is not necessarily bad for, or even present in, *both* sides. For example, it could be that, relative to objective probabilities, only plaintiff counsel is biased. Or both plaintiff and defendant counsel are biased—relative to objective probabilities—against plaintiff success, but plaintiff counsel less so because this “status quo bias” is partially offset by the role-induced bias under investigation in this paper.

ten U.S. law schools in six of the seven semesters from spring 2016 through spring 2019. Most students in these classes are second-year law students, some are third-year law students or foreign masters (LL.M.) students, a few spring students may be first-year law students, and rarely there is a cross-registrant. The exercise takes place in the last third of the semester in the context of teaching the law against insider trading. The purpose of the exercise is to engage students in a cutting-edge topic of insider trading law, while hopefully demonstrating to students the existence of own-side bias through a post-survey review of their debriefing survey answers (a hope that was never disappointed, although the instructor took pains not to give any hint before and during the exercise that this is what the instructor expected).

The basic setup is the same in all semesters and for all participants. As a mandatory classroom activity, students conduct a mock Supreme Court oral argument of *U.S. v. Newman*, the most important insider trading case of the last two decades.³ In groups of three, one student assumes the role of counsel for petitioner (the United States), one for respondent (Newman), and one the role of justice. The time allotted to the argument was 24 minutes in spring 2016 and 10 minutes in the other semesters, divided into opening arguments (10 or 4 minutes, respectively) and rebuttals (2 or 1 minute, respectively) by both sides. These timing rules were communicated to students orally and on a slide shown in the classroom.

After completion of the oral argument, students were directed to complete the anonymous debriefing survey available at a link shown on a new slide. The survey first asked students who they predicted would win at the Supreme Court, and how certain they were of their prediction on a scale “from 50% (50/50 chance, not at all certain) to 100% (completely certain).” On the next three screens, participants were asked about their assigned role, their petitioner-respondent-judge group number, and gender (female/male/prefer not to say), respectively.⁴ Finally, the last screen asked if the student permitted use of their answer for research purposes. The full text of the survey is in Online Appendix 3.

In spring 2016, students received the materials and their role assignment five days prior to the exercise, before they had even read the foundational insider trading cases preceding *U.S. v. Newman*, which were the other reading assignment for the day of the exercise. The spring 2016 materials included an edited version of the Circuit Court’s opinion in *U.S. v. Newman* of about 6,400 words. In the other semesters,

³ *Supra* note 1. In that case, two portfolio managers at hedge funds, Newman and Chiasson, had traded on insider information from two companies’ finance departments. However, the defendants did not receive the information directly from the companies’ insiders (the “tippers”) but through a chain of “tippees.” In the earlier case of *Dirks v. S.E.C.*, 463 U.S. 646 (1983), the Supreme Court had held that the tippee is criminally liable only if the tipper breaches a duty by disclosing the information and, on peculiar facts, that such a breach occurs only if the tipper received a personal benefit from the tip. Newman and Chiasson argued that the tippers in their case had not received a benefit qualifying under the *Dirks* test, and that even if the tipper had received such a benefit, Newman and Chiasson had not known about it, as would be required for a criminal conviction. The district court, however, instructed the jury that a conviction merely required a breach by the tipper with knowledge of the tippee. After Newman and Chiasson had been found guilty by the jury, they appealed to the Second Circuit, which reversed on the grounds that a conviction required a potential gain of a pecuniary or similarly valuable nature for the tippee.

⁴ Technically, the combination of group number and role identified the student to anyone who remembered which students were in which group number (which was always theoretically observable to anyone in the classroom) and who played which role (which was observable to every member in the group, theoretically to anyone in the classroom in round 3-6, and to the instructor in round 1). Realistically, however, only the students in the group would retain the number and only for the duration of the exercise, if at all (cf. discussion of group clustering and fixed effects in section 6).

students received the materials and their role assignments only in class after having read and discussed the basic law of insider trading. These materials were only 1,169 words, and students were given 15 minutes (spring 2017, spring 2018), 20 minutes (fall 2016), or 30 minutes (fall 2018, spring 2019) to read before the argument. In spring and fall of 2016, students were informed that the Supreme Court did not grant the government's petition for certiorari, but that it did grant certiorari in another case that presented the same legal question and that would soon be (spring 2016), or recently was (fall 2016), heard on oral argument. In the later semesters, no information about the state of the case was given. In all semesters, the materials were distributed in paper form. The study materials for spring 2016 and later semesters are reproduced in Online Appendices 1 and 2, respectively.

4.2 Method

4.2.1 Randomization (Main Hypotheses)

The method for testing the main hypotheses is by random assignment to the roles of petitioner's counsel, respondent's counsel, or justice (main hypothesis 1), and, in fall 2018 and spring 2019, random display of a debiasing message to half the participants at the beginning of the debriefing survey (main hypothesis 2).

4.2.1.1 Role Assignment (Main Hypothesis 1)

In spring 2016, (1) the role randomization was performed by means of Stata's random number function and emailed to students five days before class, and (2) groups were formed in class by drawing group numbers from a hat. In the other semesters, the role randomization and group composition was performed in class just before the exercise. In fall 2016, roles and groups were drawn from a hat. In the remaining four semesters, (1) groups were formed by dividing the classroom into three equally sized portions, assigning numbers 1 through $N/3$ to students in each portion (following rows of seats either in switchback or parallel form), and matching students with the same number from the three divisions, and (2) roles were assigned by alphabetical order of last name within group. In each case, unmatched residual students were asked to team up with another student with the same role.⁵

The assignment and recording of roles raise four minor concerns of potential bias. First, roles were self-reported. If students reported their role with error, treatment would be mismeasured, biasing effect size estimates downwards. However, students presumably had a fresh clear memory of their role such that misreporting should be rare, and effect size estimates are large anyway. Second, students could have assumed a role different from the one assigned. If they did so because they believed in the substantive position of their assumed role, it would bias estimated treatment effects upwards. However, only one of 75 students reported an actual role different from the assigned role in spring 2016, the only semester when the survey asked about both.⁶ Moreover, students would have had to "collude" with the other two students in their group to change their role in defiance of the instructor's instructions, which seems a priori very unlikely. Third, in the four semesters when roles within each group were assigned by last name, role could be correlated with ethnic or national origin because of the latter's links to certain names. However, there is no reason to think origin is correlated with predictions regarding the *Newman* case, and

⁵ In spring 2019, one student and the student's group partner (the group was incomplete) were excluded from the analysis because that student had just heard about the *Newman* case in another class taught by the instructor.

⁶Results are unchanged if this student is counted in the assigned role (as reported here), in the actual role, or dropped.

in any event the results hold even when restricting the sample to spring and fall 2016, when last name played no role in the randomization. Fourth and finally, the creation of groups from different sections of the classroom may have affected the distribution of combinations of student types (e.g., if sitting in the left/middle/right of the classroom corresponded to political preferences, groups were essentially guaranteed to be ideologically mixed). However, this distribution is orthogonal to the treatment assignment and hence cannot create bias.

4.2.1.2 *Debiasing Message (Main Hypothesis 2)*

The debiasing message was randomly shown to only half the participants in the fall of 2018 and spring 2019 using Qualtrics's randomization feature (in other semesters, no debiasing message was shown). This randomization was independent of the role randomization. When shown, the debiasing message was inserted after the survey's opening line

"Please answer a couple of anonymous questions about U.S. v. Newman, which you have just argued/heard."

and read:

"If your role was one of petitioner or respondent: Please bear in mind that your role was that of an advocate, whereas we are now stepping back out and I am asking your opinion as a neutral, sober observer."

Thus, the debiasing intervention consisted of a simple message that emphasized the difference between advocacy and neutral assessment. Section 7 will discuss the possibility of stronger debiasing.

4.2.2 *Effect Size Differentials by Group (Ancillary Hypotheses)*

The relative effect size in spring 2016 as compared to other semesters and between male and female participants sheds light on ancillary hypotheses 1 and 2, respectively. It is important to note, however, that role exposure time and gender are not randomly assigned and hence potentially confounded with correlated differences between these groups.

Role exposure time (spring 2016 vs. other semesters) is confounded with any random difference of the instructor's teaching in spring 2016 relative to other semesters. Moreover, only in that semester did students receive a multi-page excerpt from the Court of Appeals opinion as part of the materials, which could well have reduced the dispersion of students' views (but, as will be seen, did not).⁷

Gender is an immutable trait, such that even the question of its causal effect is arguably ill-defined. In any event, gender might be correlated with application, admission, and enrollment patterns. That said, the comparison of male and female participants can at least shed light on any differences in these groups conditional on application, admission, and enrollment.

5 Data

A total of 391 students participated and gave permission to use their data. Owing to the nature of the data collection, sample size was not determined ahead of time. It was decided not to collect more data


⁷ The failure of the lower court's opinion to change students' views might be a manifestation of the phenomenon that people tend to question legal experts' authority rather than change their views of the law if the two conflict (Simon & Scurich 2013).

after spring 2019 because the instructor would not teach the same class for at least two years and because *U.S. v. Newman* was getting stale. There was also an increasing risk that students would know about the exercise from students in prior years, although in general that risk should not be higher than in other lab studies with sequential participation and the risk was mitigated by spreading the study over two schools. The pattern of data over time allays concerns that results are due to contagion or *p*-hacking. The basic effect was observed in each of the six rounds (cf. figure 2), was statistically significant at the 5% (10%) level in each of rounds 1 and 4 (3 and 5) even considered in isolation, and the *p*-value using all then available data was strictly decreasing over the course of the six rounds.

In each of the six semesters, the total number of participants (75, 58, 32, 76, 62, and 91, respectively) was between 9% and 28% lower than official enrollment (94, 64, 43, 93, 74, 104). The reason is presumably that students were absent from class for various reasons. They might also have been present but refused to fill in the exit survey, but that should be rare given its mandatory nature. Two students filled in the survey but did not give permission to use their data for research purposes, and one started the survey but did not submit a prediction.

Imbalance in treatment group sizes could arise from the post-randomization attrition just discussed or from error in self-reported treatment status discussed in section 4.2.1.1. As shown in figure 1, however, imbalances are very minor. The petitioner, justice, and respondent group sizes differ only slightly more than the randomization mechanism’s maximum of one in spring 2016 (2), spring 2017 (2), spring 2018 (3), and fall 2018 (2).

Figure 1: Randomization Balance

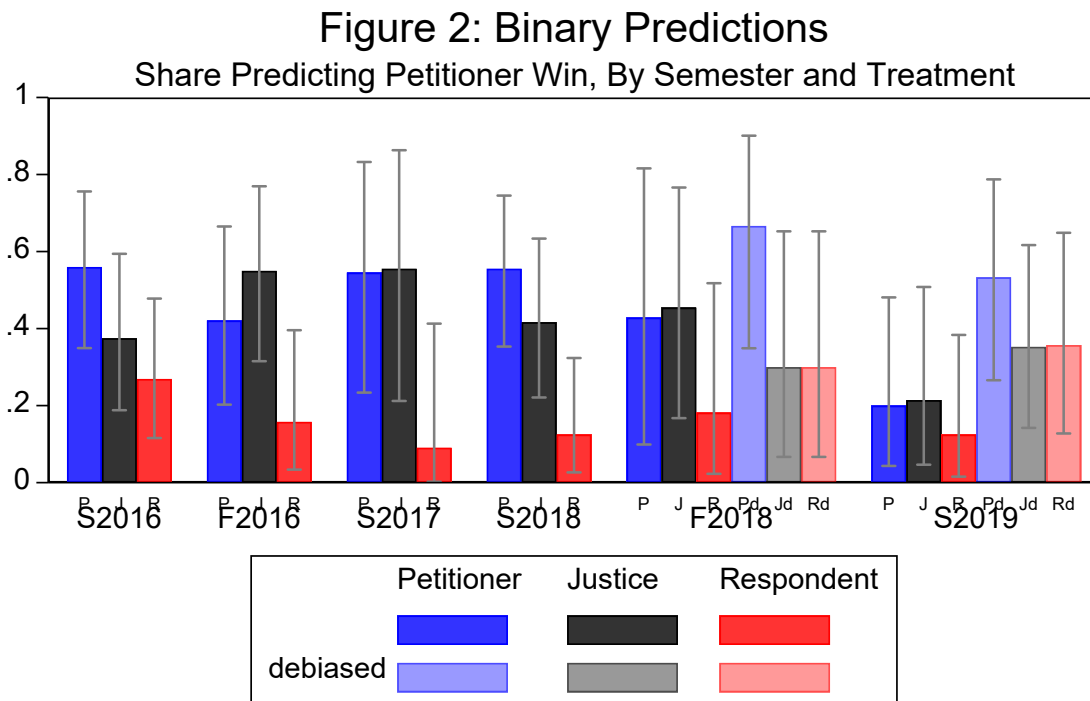


For each semester (Spring/Fall-Year), the graph plots the number of complete participations in each of the three treatment levels (Petitioner, Justice, Respondent) broken down by gender.

Figure 1 also shows that the randomization was successful in achieving covariate balance at least with respect to the only observable covariate, gender. In total, there were 189 female participants, 198 male participants, and 4 who “prefer[red] not to say” their gender. They were randomized into the three treatment levels within each semester, and for all semesters combined, in roughly equal proportions except in Spring 2017. A Fisher exact test rejects the null of equal distribution at the 10% level only in Spring 2017 ($p=0.03$) and not for all semesters combined.

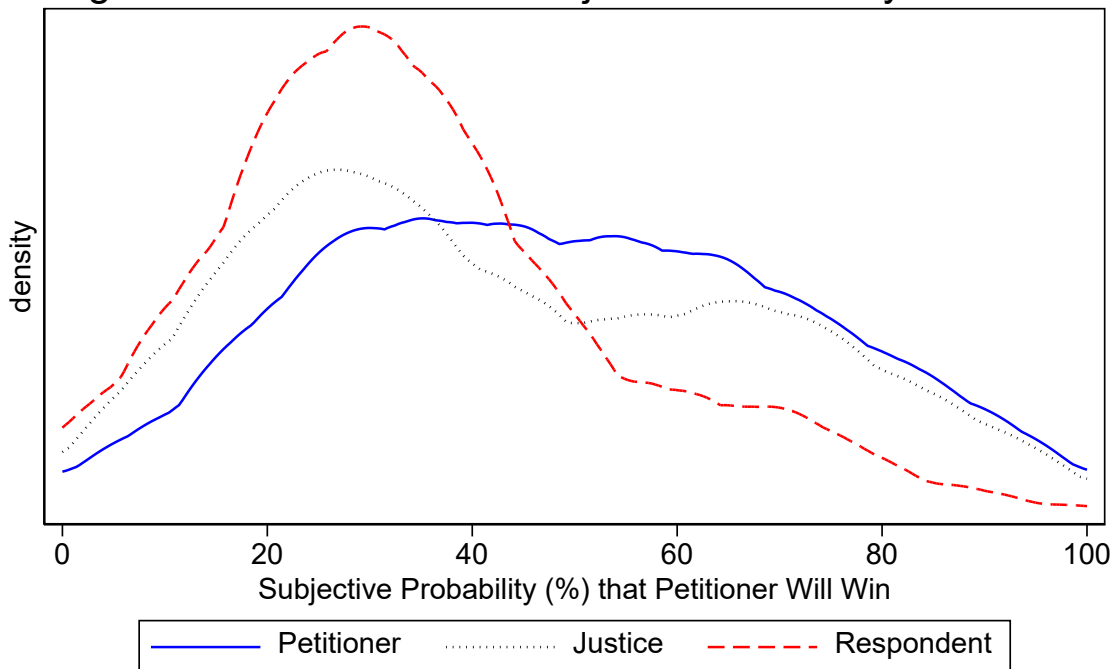
6 Results

Figures 2 and 3 present the raw outcome data. Figure 2 shows the rate of participants who predicted the petitioner to win separately for each semester and treatment condition. The difference between petitioner’s and respondent’s counsel roles is readily apparent, regardless of debiasing. In each semester and with or without debiasing, students representing petitioner were more likely to predict the petitioner to win, usually by a large margin. The same picture emerges from figure 3’s role-specific kernel density plots of participants’ subjective probability estimates that petitioner will win. The subjective probability estimates are constructed as the percentage certainty for those who predicted petitioner to win, and 100 minus the percentage certainty for those who predicted the respondent to win. The belief distributions are neatly ordered in terms of first-order stochastic dominance: petitioner counsel’s distribution dominates justice’s distribution, which in turn dominates respondent counsel’s distribution.



For each semester (Spring/Fall-Year), the graph plots the share and exact 95% c.i. of participants predicting petitioner win by main treatment level (Petitioner, Justice, Respondent) and, where applicable, by debiasing (Fall 2018 and Spring 2019).

Figure 3: Distributions of Subjective Probability Estimates




Epanechnikov kernel density plots of subjective certainty that petitioner will win (on a scale of 50%-100%), or 100 minus the subjective certainty for those who predicted petitioner will lose.

To be sure, the wide 95% confidence intervals in figure 2 caution against overinterpreting individual bars. For example, justices predict petitioner win at even higher rates than petitioner’s counsel in Fall 2016, Spring 2017, and, when not debiased, Fall 2018 and Spring 2019. But this is probably just random noise: on average, justices predict petitioner win at a rate (0.40) that lies between petitioner counsel’s (0.50) and respondent counsel’s (0.20) and that is not statistically distinguishable from the midpoint between the latter two (Fisher’s $p=0.32$). Nevertheless, the cumulative evidence for role-induced bias is strong. The next two figures and associated discussion examine this formally.

Figure 4 begins with simple univariate tests. It shows the treatment effect corresponding to main hypothesis 1 as the simple difference between the rates at which petitioner’s and respondent’s counsel predicted petitioner to win. As the next paragraph will explain, this perspective is not fully statistically accurate, but it has the advantage of being simple and transparent. Across all treatment conditions, the estimated effect size is 0.30 ± 0.11 , which is highly significantly different from zero (Fisher’s two-sided $p < 10^{-6}$, or $p < 10^{-5}$ after Bonferroni-adjustment for this paper’s multi-testing of two main and two ancillary hypotheses). There are only minor differences in estimated effect sizes between those who received their role assignment pre-class (i.e., in spring 2016) (0.29 ± 0.26) and those who received it in class, and within that latter group between those who received (0.26 ± 0.26) and those who did not receive (0.31 ± 0.13) the debiasing message (the confidence intervals vary in width due to the differences in sample size). The difference in estimated effect sizes between male or “rather not say” (0.35 ± 0.16) and female (0.26 ± 0.15) participants is a little larger. However, none of these differences in estimated effect sizes is even close to statistically significant, as suggested by the wide univariate confidence intervals and confirmed by the multivariate analysis below.

Figure 4: Main Treatment Effect (Univariate)



The foregoing univariate analysis failed to deal with two types of statistical complications. First, the estimates did not account for the facts that the data were drawn from six different semesters with possibly different means, and that the assignment of one of the two randomized treatments (role, debiasing) might not have been completely balanced across that of the other. While the two treatments are independent in expectation, slight imbalances did occur by random chance. Second, the statistical inference did not account for the corresponding loss of degrees of freedom and, more importantly, the fact that individual observations are not independent within group because group members potentially influenced each other's views. Modelling the participants' predictions that the petitioner will win can address these issues. The estimated linear probability model is

$$prediction_{igs} = \alpha_{gs} + \beta_r d_{i,r} + \sum_{c \in \{D,F\}} \gamma_c d_{i,c} + \sum_{c \in \{D,PC,F\}} \delta_{c,r} d_{i,c} d_{i,r} + \varepsilon_{igs} \quad (1)$$

where $prediction_{igs}$ is the prediction of participant i in group g in semester s (binary as in figures 2 and 5, or the subjective probability estimate as in figure 3); α_{gs} is a fixed effect for group g in semester s (constrained to be constant for all groups g in semester s in all but the fixed effect specifications reported below); $d_{i,r}$ is a dummy indicating that participant i was assigned role $r \in \{P, J\}$ (petitioner or justice, respondent being the omitted category); $d_{i,c}$ is a dummy indicating that participant i was assigned, or belonged, to condition $c \in \{D, PC, F\}$ (debiasing, pre-class role assignment [spring 2016], and female

gender, respectively); and ε_{igs} is an individual error term.⁸ In this specification, β_P identifies the main effect of the role assignment (i.e., the role-induced bias), and $\delta_{c,P}$ identifies the interaction effect with (i.e., the attenuation or accentuation of the bias by) condition c .

Clustering the standard errors by group gs allows for arbitrary correlation of errors ε_{igs} within group. However, there is some concern about the reliability of the group membership information that survey participants entered 30-40 minutes after receipt and use of this subsequently useless number. Under the study design, at most two roles of petitioner/respondent/justice could be filled at most double in at most one group per session. But the data from the six sessions contain two triple-filled roles and fifteen double-filled roles.⁹ Five participants did not enter a group number at all. Fortunately, as can be seen in figure 5, it turns out that the estimates are essentially identical without clustering, with clustering on group using only “clean” groups (i.e., excluding groups with double-filling), with clustering on group using all groups, or adding group fixed effects.

Figure 5: Treatment Effects and Interactions (Multivariate)

OLS estimates and 95% c.i. of equation (1), dep. var.: binary prediction




Figure 5 shows the main coefficients of interest with 95% confidence intervals from OLS estimates of equation (1) with the binary prediction as the dependent variable. (Using percentage predictions as

⁸ Note that pre-class role assignment PC is constant within each semester, such that only its interaction with treatment $d_{i,PC}d_{i,r}$ but not its level $d_{i,PC}$ is identified and thus included in the estimated equation.

⁹ Alternatively, participants might not have followed the design instructions and worked in teams. However, the instructor did not observe such behavior in the classroom.

dependent variables yields a very similar picture with larger confidence intervals that still do not include zero for the main effect.) The multiple estimates per coefficient correspond to four different variants of the equation (semester or group fixed effects, gender controls or not) and the three aforementioned different ways of calculating the standard errors (which lead to slightly different sample sizes and hence also affect the point estimates). As can be seen, $\beta_p \approx 0.30$ in all specifications, with a lower 95% confidence bound far above zero. By contrast, the interaction effects are not merely not statistically significant but estimated to be close to zero or even in the opposite of the expected direction. The highest estimate of a debiasing effect is $\delta_{D,p} = -0.08$. The interactions of pre-class exposure ($\delta_{PC,p}$) and female gender ($\delta_{F,p}$) with the bias are either estimated to be exactly zero, or to go slightly in the unexpected direction. To be sure, the confidence intervals for the interaction effects are wide, so that estimates in the unexpected direction are presumably noise and even sizeable effects in the hypothesized direction cannot be ruled out. That said, the estimates certainly do not provide evidence for the hypothesized effects.

As in Eigen & Listokin (2012), there is no gender effect. Women also reported almost exactly as much subjective confidence as men on average: 70.6% certain on a scale of 50%-100%, barely shy of men's 72.5%. A t -test does not reject equality of the two means ($p=0.15$) in spite of the fairly large sample size (368; this is slightly smaller than the sample for the main effects because not all participants submitted a confidence rating).

Interestingly, students who received the debiasing message were more likely to predict petitioner win regardless of their role assignment. This can be seen in figure 2 and is confirmed by a simple univariate Fisher test ($p=0.03$) and equation 1's estimated coefficient—using the binary prediction as dependent variable—of $\hat{\gamma}_D \approx 0.20 \pm 0.21$ (although the estimate is reversed for students in the justice role since $\hat{\delta}_{D,J} \approx -0.30$). A possible explanation is that the reminder to remain neutral pushed students' predictions towards equipoise, whereas they otherwise leaned against the petitioner's position. Across all roles, students in the debiasing condition predicted petitioner win at a rate of 0.42, which is not statistically different from 0.50 ($p=0.21$ in an exact binomial test). By contrast, students in the baseline predicted petitioner win at a rate of only 0.35 across all semesters and even only 0.24 in Fall 2018 and Spring 2019 when the debiasing message was randomly shown; both are statistically different from 0.50 at $p \leq 10^{-5}$. That said, if the debiasing message pushed towards equipoise, one might also expect predictions to be less confident, but there is no evidence for that.¹⁰ This is an interesting question for future research.

7 Discussion

The present experiment has shown that law students exhibit substantial role-induced bias with respect to law-based predictions—30%—within 40 minutes of randomized role assignment. Eigen & Listokin (2012) had found similar bias only after much longer engagement with the case (over 30 hours), whereas the other literature reviewed in section 2 had found role-induced bias similarly fast but in contexts further removed from law (laypeople, factual questions). That the bias arises so fast, and that legal training does not inoculate against it, is an important, if concerning, finding for the legal profession and its ability to provide advice to its clients. It is even more concerning because the experiment also showed that a simple

¹⁰ Students in the debiasing condition are actually one percentage point more certain in their prediction than those in the baseline. Controlling for covariates (by estimating equation 1 with self-reported confidence as the dependent variable) flips the sign of the one percentage point difference but this estimate is so imprecise as to be useless, with a standard error about four times the absolute point estimate.

reminder of the need for neutrality, similar perhaps to what lawyers might attempt in practice, does not attenuate the bias, at least not meaningfully.

An important feature of the design was that students first argued the case and then made their predictions. It is likely that engaging in the argument lead to information distortions that persisted even after the end of the argument (cf. DeKay 2015, Simon et al. 2015, Russo 2018). Had the students approached the case in a more neutral manner first before being assigned their roles, they might have formed more neutral expectations, which might have persisted. The order of the experiment appears more realistic, however, because lawyers tend to learn about a case only after being engaged for one side by a client. That said, the experiment was unrealistic in as much as real-world lawyers make predictions before they engage in litigious argument. Exploring such order effects in future research is both practically and theoretically important.

More generally, the present findings raise the question what, if anything, could eliminate or even reduce the bias. Debiasing is generally known to be difficult (e.g., Russo 2018). While a simple reminder as in the present experiment had been found ineffectual in other settings, weighing arguments had been found effective (Simon 2004). Lawyers, of course, are trained to, and do, weigh arguments all the time – but this did not inoculate the experiment participants against their bias. That said, some more elaborate debiasing tasks have been found effective. Elimination of role-induced bias was achieved in Babcock et al. (1997) by asking participants to write down the weakness of their case, and in Schweizer (2018) by asking participants to submit separately probability estimates for the ingredients of a Bayesian calculation and showing them the result of that calculation; both of these studies used non-lawyers and non-legal questions. A related question is if some individuals are less prone to bias than others. For example, the present experiment's finding of a 30% difference in average prediction is consistent with 70% participants being unbiased. From both a theoretical and a practical point of view, these are important questions for future research.

References

- Babcock, Linda, George Lowenstein, and Samuel Issacharoff. 1997. Creating Convergence: Debiasing Biased Litigants. *Law & Social Inquiry* 22:913-925.
- Babcock, Linda, George Lowenstein, Samuel Issacharoff, and Colin Camerer. 1995. Biased Judgments of Fairness in Bargaining. *American Economic Review* 85:1337-1343.
- Braman, Eileen, and Thomas E. Nelson. 2007. Mechanism of Motivated Reasoning? Analogical Perception in Discrimination Disputes. *American Journal of Political Science* 51:940–956.
- DeKay, Michael L. 2015. Predecisional Information Distortion and the Self-Fulfilling Prophecy of Early Preferences in Choice. *Current Directions in Psychological Science* 24:405–11.
- Eigen, Zev J., and Yair Listokin. 2012. Do Lawyers Really Believe Their Own Hype, and Should They? A Natural Experiment. *Journal of Legal Studies* 41:239-267.
- Engel, Christoph, and Andreas Glöckner. 2013. Role-Induced Bias in Court: An Experimental Analysis. *Journal of Behavioral Decision Making* 26: 272–284 (2013).
- European Parliament (Directorate General of Internal Policies of the Union, Policy Department for Citizens' Rights and Constitutional Affairs). 2017. Mapping the Representation of Women and Men in Legal Professions Across the EU. Study PE 596.804, August 2017.
- Furgeson, Joshua R., Linda Babcock, and Peter M. Shane. 2008a. Do a Law's Policy Implications Affect Beliefs About Its Constitutionality? An Experimental Test. *Law and Human Behavior* 32:219-227.
- Furgeson, Joshua R., Linda Babcock, and Peter M. Shane. 2008b. Behind the Mask of Method: Political Orientation and Constitutional Interpretive Preferences. *Law and Human Behavior* 32:502-510.
- Goodman-Delahunty, Jane, Pär Anders Granhag, Maria Hartwig, and Elizabeth F. Loftus. 2010. Insightful or Wishful: Lawyers' Ability to Predict Case Outcomes. *Psychology, Public Policy, and Law* 16:133–157.
- Hippel, Svenja, and Sven Hoeppe. 2019. Biased judgements of fairness in bargaining: A replication in the laboratory. *International Review of Law and Economics* 58:63-74.
- Loewenstein, George, Samuel Issacharoff, Colin Camerer, and Linda Babcock. 1993. Self-Serving Assessments of Fairness and Pretrial Bargaining. *Journal of Legal Studies* 22:135-159.
- Moore, Don A., and Amelia S. Dev. 2018. Individual Differences in Overconfidence. In V. Zeigler-Hill and T. K. Shackelford (Eds.), *Encyclopedia of Personality and Individual Differences*. New York: Springer.
- Russo, J. Edward. 2018. Bayesian Revision vs. Information Distortion. *Frontiers in Psychology* 9:1550.
- Schweizer, Mark. 2018. De-biasing role induced bias using Bayesian networks. Available at <https://ssrn.com/abstract=3231859>.
- Sharot, Tali. 2011. The Optimism Bias. *Current Biology* 21:R941-R945.
- Simon, Dan. 2004. A Third View of the Black Box: Cognitive Coherence in Legal Decision Making. *University of Chicago Law Review* 71:511-586.

- Simon, Dan, Chadwick J. Snow, and Stephen J. Read. 2004. The Redux of Cognitive Consistency Theories: Evidence Judgments by Constraint Satisfaction. *Journal of Personality and Social Psychology* 86:814–837.
- Simon, Dan, and Nicholas Scurich. 2013. The Effect of Legal Expert Commentary on Lay Judgments of Judicial Decision Making. *Journal of Empirical Legal Studies* 10:797–814.
- Simon, Dan, Douglas M. Stenstrom, and Stephen J. Read. 2008. On the Objectivity of Investigations: An Experiment. Available at <https://ssrn.com/abstract=1130103>.
- Simon, Dan, Douglas M. Stenstrom, and Stephen J. Read. 2015. The Coherence Effect: Blending Cold and Hot Cognitions. *Journal of Personality and Social Psychology* 109:369–394.

Online Appendix 1: Spring 2016 Study Materials

In the prior class five days before the study, a printout of the following document was distributed to students. While the document was passed around, the instructor explained orally that students should prepare to argue this case in the next week in the role of petitioner, respondent, or justice, as instructed by an email that an automated system sent to them as the instructor was speaking (with the simple text: “You have been assigned the role of ____.”).

“US v. Newman

“For our class on Monday, 3/28, please come prepared to argue your assigned side (government-petitioner, defendant-respondent, or your favorite Justice) in a mock Supreme Court oral argument in *U.S. v. Newman*. The case is described below. You will receive your role assignment by e-mail.

To prepare for the mock argument, you will have to read not only the decision below but also the relevant Supreme Court precedents *Chiarella*, *Dirks*, and *O’Hagan*, as assigned on the syllabus and available on [course website] and in print. If you want, you can read more than the passages excerpted on [course website], but I do not expect you to.

Please note: In actuality, the Supreme Court did not grant the government’s petition for certiorari in *U.S. v. Newman*. But it recently granted certiorari in a case that presents the same legal question in a less relevant setting. That case will be heard later this term.

QUESTION PRESENTED

In *Dirks v. SEC*, 463 U.S. 646 (1983), this Court held that a corporate insider breaches his fiduciary duty and subjects himself to insider-trading liability when he personally benefits from the selective disclosure of material, nonpublic information for securities trading—including when he “makes a gift of confidential information to a trading relative or friend.” *Id.* at 664. “The tip and trade,” the Court explained, “resemble trading by the insider himself followed by a gift of profits to the recipient.” *Ibid.*

The question presented is whether the court of appeals erroneously departed from this Court’s decision in *Dirks* by holding that liability under a gifting theory requires “proof of a meaningfully close personal relationship that generates an exchange that is objective, consequential, and represents at least a potential gain of a pecuniary or similarly valuable nature.”

PARTIES TO THE PROCEEDING

Petitioner is the United States of America, which was appellee in the court of appeals. Respondents are Todd Newman and Anthony Chiasson, who were appellants in the court of appeals.”

The original instructions then included a printout of an edited version of the 2nd Circuit’s opinion in the case, which is available at <https://h2o.law.harvard.edu/collages/40640>.

Online Appendix 2: Study materials (other than for spring 2016)

“U.S. v. Newman

The question in this case is whether a conviction for "tippee" insider trading under SEC rule 10b-5 requires that the tipper received at least a potential gain of a pecuniary or similarly valuable nature in exchange for the leaked information.

The Government brought charges against Todd Newman, a portfolio manager at Diamondback Capital Management, LLC, and Anthony Chiasson, a portfolio manager at Level Global Investors, L.P. At the six-week jury trial, the Government presented evidence that a group of financial analysts exchanged information they obtained from company insiders, both directly and more often indirectly. Specifically, the Government alleged that these analysts received information from insiders at Dell and NVIDIA disclosing those companies' earnings numbers before they were publicly released in Dell's May 2008 and August 2008 earnings announcements and NVIDIA's May 2008 earnings announcement. These analysts then passed the inside information to their portfolio managers, including Newman and Chiasson, who, in turn, executed trades in Dell and NVIDIA stock, earning approximately \$4 million and \$68 million, respectively, in profits for their respective funds.

Newman and Chiasson were several steps removed from the corporate insiders and there was no evidence that either was aware of the source of the inside information:

- With respect to the Dell tipping chain, the evidence established that Rob Ray of Dell's investor relations department tipped information regarding Dell's consolidated earnings numbers to Sandy Goyal, an analyst at Neuberger Berman. Goyal in turn gave the information to Diamondback analyst Jesse Tortora. Tortora in turn relayed the information to his manager Newman as well as to other analysts including Level Global analyst Spyridon "Sam" Adondakis. Adondakis then passed along the Dell information to Chiasson, making Newman and Chiasson three and four levels removed from the inside tipper, respectively.
- With respect to the NVIDIA tipping chain, the evidence established that Chris Choi of NVIDIA's finance unit tipped inside information to Hyung Lim, a former executive at technology companies Broadcom Corp. and Altera Corp., whom Choi knew from church. Lim passed the information to co-defendant Danny Kuo, an analyst at Whittier Trust. Kuo circulated the information to the group of analyst friends, including Tortora and Adondakis, who in turn gave the information to Newman and Chiasson, making Newman and Chiasson four levels removed from the inside tippers.

The Government charged that Newman and Chiasson were criminally liable for insider trading because, as sophisticated traders, they must have known that information was disclosed by insiders in breach of a fiduciary duty, and not for any legitimate corporate purpose.

At the close of evidence, Newman and Chiasson moved for a judgment of acquittal pursuant to Federal Rule of Criminal Procedure 29. They argued that there was no evidence that the corporate insiders provided inside information in exchange for a personal benefit which is required to establish tipper liability under *Dirks v. S.E.C.*, 463 U.S. 646 (1983). Because a tippee's liability derives from the liability of the tipper, Newman and Chiasson argued that they could not be found guilty of insider trading. Newman and Chiasson also argued that, even if the corporate insiders

had received a personal benefit in exchange for the inside information, there was no evidence that they knew about any such benefit. Absent such knowledge, appellants argued, they were not aware of, or participants in, the tippers' fraudulent breaches of fiduciary duties to Dell or NVIDIA, and could not be convicted of insider trading under *Dirks*. In the alternative, appellants requested that the court instruct the jury that it must find that Newman and Chiasson knew that the corporate insiders had disclosed confidential information for personal benefit in order to find them guilty.

The district court denied the appellants' Rule 29 motions and did not give Newman and Chiasson's proposed jury instruction. The district court relied on language in *Dirks* (at 659) that "a tippee assumes a fiduciary duty to the shareholders of a corporation not to trade on material nonpublic information ... when the insider has breached his fiduciary duty to the shareholders by disclosing the information to the tippee and the tippee knows or should know that there has been a breach." Holding that a tipper could breach a duty to the corporation even if the tipper did not receive a personal benefit, the district court gave the following instructions on the tippers' intent and the personal benefit requirement:

Now, if you find that Mr. Ray and/or Mr. Choi had a fiduciary or other relationship of trust and confidence with their employers, then you must next consider whether the Government has proven beyond a reasonable doubt that they intentionally breached that duty of trust and confidence by disclosing material, nonpublic information for their own benefit.

On the issue of the appellants' knowledge, the district court instructed the jury:

To meet its burden, the Government must also prove beyond a reasonable doubt that the defendant you are considering knew that the material, nonpublic information had been disclosed by the insider in breach of a duty of trust and confidence. The mere receipt of material, nonpublic information by a defendant, and even trading on that information, is not sufficient; he must have known that it was originally disclosed by the insider in violation of a duty of confidentiality.

The jury returned a verdict of guilty on all counts. The district court sentenced Newman and Chiasson to prison terms of 4.5 and 6.5 years, respectively.

On appeal, the U.S. Court of Appeals for the Second Circuit vacated the convictions and remanded for the district court to dismiss the indictment with prejudice for insufficiency of evidence. The Court of Appeals held that in order to sustain a conviction for insider trading, the Government must prove beyond a reasonable doubt that the tippee knew that an insider disclosed confidential information *and* that he did so in exchange for a personal benefit. Moreover, the Court of Appeals held that the existence of a benefit could not merely be assumed. Rather, the Court of Appeals held that insider trading liability under a gifting theory requires proof of a meaningfully close personal relationship that generates an exchange that is objective, consequential, and represents at least a potential gain of a pecuniary or similarly valuable nature. In support of its position, the Court of Appeals relied on *Dirks*, which had held (at 661-662): "All disclosures of confidential corporate information are not inconsistent with the duty insiders owe to shareholders. ... Whether disclosure is a breach of duty therefore depends in large part on the purpose of the disclosure. ... [T]he test is whether the insider personally will benefit, directly or indirectly, from his disclosure. Absent some personal gain, there has been no breach of duty to stockholders. And absent a breach by the insider, there is no derivative breach."

The Government petitioned the Supreme Court for a writ of certioari. It argued that the Court of Appeals misapplied *Dirks*. According to the Government, the facts in *Dirks* are distinguishable, and the benefit test of mere evidentiary nature specific to the facts in *Dirks*.”

Only in fall 2016:

[“The Supreme Court granted cert and the case is scheduled for oral argument this term (fall 2016).¹¹”]

¹¹ Actually, the Supreme Court granted cert in a different case raising the same legal question, and oral argument was held a month ago. But pretend this is the real case and oral argument is yet to come, with you doing the argument or the questioning.

Online Appendix 3: Survey Instrument

Screen 1:

Please answer a couple of anonymous questions about U.S. v. Newman, which you have just argued/heard.

[If your role was one of petitioner or respondent: Please bear in mind that your role was that of an advocate, whereas we are now stepping back out and I am asking your opinion as a neutral, sober observer.]¹²

How do you predict the Supreme Court will decide? Who is more likely to win? [Petitioner (Government)/Respondents (Newman and Chiasson)]

How certain are you of your prediction, from 50% (50/50 chance, not at all certain) to 100% (completely certain)? [slider with choices 50-100]

Screen 2:

[Spring 2016:] What was your assigned role, as per the email?
[Petitioner/Respondent/Justice]

Did you play your assigned role in the mock argument? [Yes/No]

[If answer to last question was “No”:] What was your actual role in the mock argument? [Petitioner/Respondent/Justice]

[All other semesters:] What was your assigned role? [Petitioner/Respondent/Justice]

Screen 3:

What was your petitioner-respondent-judge group? Please write the group number.

Screen 4:

What is your gender? [Female/Male/Prefer not to say]

Screen 5:

Last question: would you allow me to use your anonymous answers for research purposes? (The answer to this question is completely anonymous as well, and you are of course completely free to answer no.) [Yes/No]

¹² Randomly shown to 50% of participants only in fall 2018 and spring 2019.