Online Appendix for “The Substantive Effects of Descriptive Representation: Gay and Lesbian Members of Congress are more Supportive of Gay Rights.”

Appendix A: Matching Overview

To illustrate the basic idea of CEM, consider the scenario depicted in Table A1 in which we have seven legislators, and we are matching LGB and non-LGB legislators-districts based on three continuous variables X1, X2, and X3. Each variable is divided into 3 bins such that legislators with similar scores on a variable are assigned to the same bin for that variable. Each combination of bins is referred to as a strata, and legislators with similar scores on all variables, as reflected by being assigned to the same combination of bins, have the same strata. Legislators that share the same strata with others are matched for purposes of comparison.

Table A1: An Example of Coarsened Exact Matching

Legislator	*LGB*	*X1bin*	*X2bin*	*X3bin*	*Strata*	*Matched*
1	1	2	3	1	1	1
2	0	3	3	3	2	0
3	0	2	3	1	1	1
4	0	1	1	1	4	0
5	0	2	3	1	1	1
6	0	1	2	3	5	0
7	0	3	2	1	6	0

In Table 1, we present a hypothetical example. Here, legislator 1 is assigned to bin 2 for X1, bin 3 for X2, and bin 1 for X3 based on their values for each of these three variables. Given these values, Legislator 1 is assigned to strata 1. As legislators 3 and 5 also share membership in these bins, they are assigned strata 1 as well. These legislators are then matched to each other for purposes of comparison. We note that no other legislators share identical combinations of bin memberships and hence there are no other legislators with matching strata.

To estimate the relationship between the outcome (i.e. HRC score) and LGB status we would then regress the HRC score on the value of X1, X2, and X3 but only for

the matched cases and we would only interpret the LGB status variable as having a causal effect on the HRC score, because the design is not suited to drawing inferences about the X variables used for matching.

Appendix B. Exclusion of Barney Frank from Analysis

To ensure our results are not the result of strategic manipulation by an influential legislator, Barney Frank. To identify if the exclusion of Barney Frank affects the results we exclude him from the analyses in Figures B1 and B2. Figure B1 reports the results of our matching-based analyses using the full complement of Congresses, but matching on a smaller subset of data because the LGB Household share is not available prior to the 109th Congress.

Figure B1: Estimated Relationship Between LGB Status and HRC Scores, excluding Democrat Barney Frank.

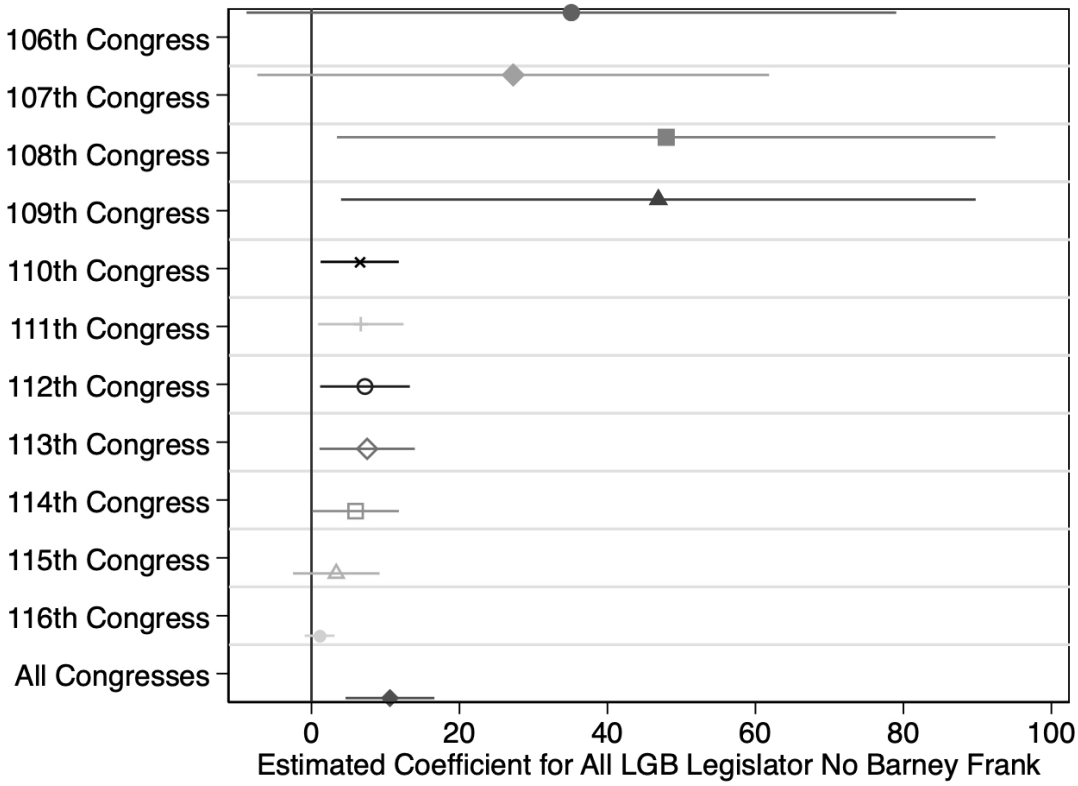

Figure B2: Estimated Relationship Between LGB Status and HRC Scores, excluding Democrat Barney Frank and using LGB Household Share

The figure below uses only the 109th and following Congresses because that is when the LGB Household variable is available.

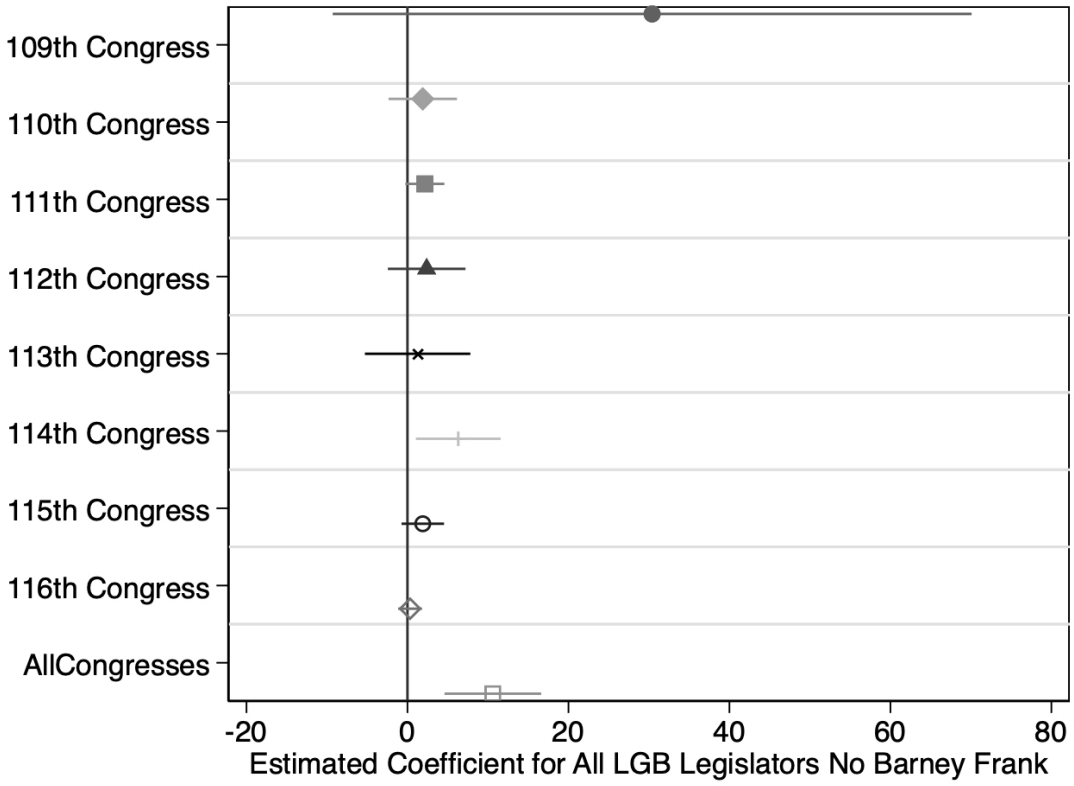

Appendix C: Results using Alternative Measures of Ideology

At the suggestion of a reviewer, we re-estimated our models using the Warshaw and Tausanovitch [1] Congressional ideology data in our analyses. Unfortunately, these data are only available for Congresses 111, 114, 116, 117 and 118. As the time-period we study covers Congresses 105-116 there are only three Congresses of overlap. This substantially reduces the number of comparisons we can make. However, as we show in Figure C1, below, using these alternative measures of public opinion based on district ideology does *not* produce substantively different results than those we present in the main body for which we used Presidential vote share to measure public opinion. More specifically, we re-estimated our analyses using all three of the alternative ideology measures and the results are essentially identical across all measures. Given the limited overlap between congresses, and the robust results we obtain, we do not include the Warshaw and Tausanovitch measure in the analyses presented in the main body of the paper.

The consistency of the results is not surprising, because these alternative ideology measures are correlated with Presidential Vote Share (how we measured ideology) at over 0.9 so our current measure is already capturing a significant amount of the variation in ideology. The other advantage of Presidential vote share is that we know it is measured correctly, whereas there is error in our measures of ideology, but this error is not incorporated into our empirical models.

[1] Christopher Warshaw; Chris Tausanovitch, 2022, "Subnational ideology and presidential vote estimates (v2022)", https://doi.org/10.7910/DVN/BQKU4M, Harvard Dataverse, V1,

Figure C1: Estimated effect of LGB status using alternative ideology measure

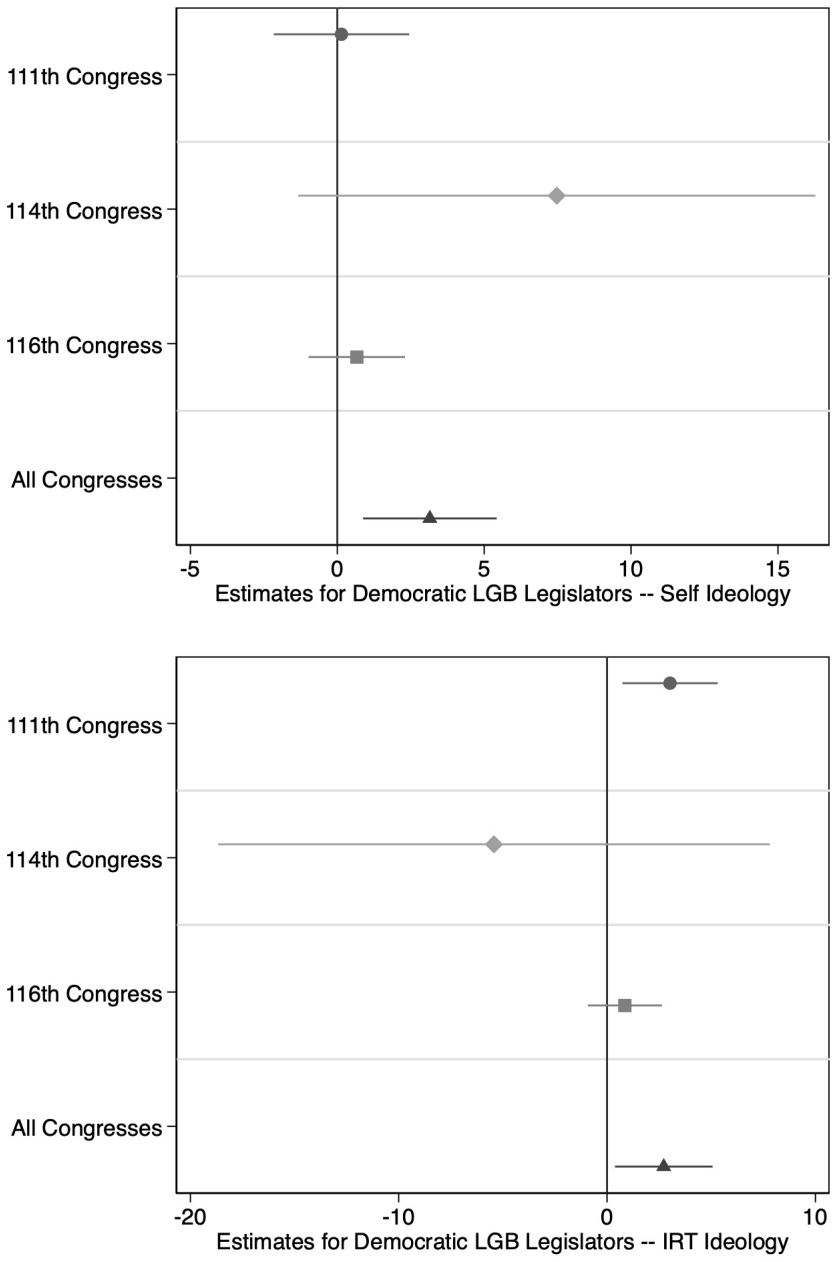

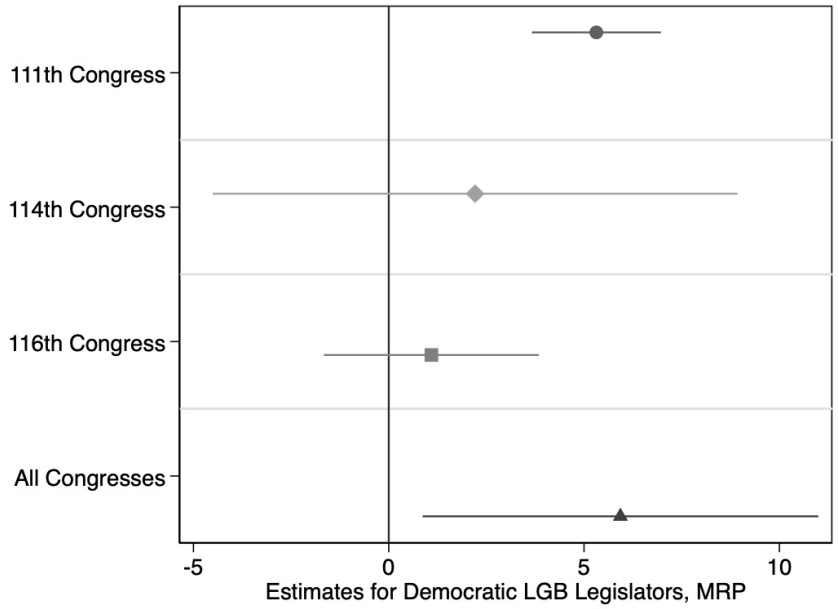
111th Congress
114th Congress
116th Congress
All Congresses
-5
0
5
10
Estimates for Democratic LGB Legislators, MRP

Appendix D: Results using Nearest Neighbor Matching

To examine the degree to which our results are robust to other matching methods, we also employ nearest neighbor matching. Specifically, we used the same variables as we did with CEM but instead employ a version of nearest neighbor matching.[2] We combined all congressional sessions together, because there are insufficient observations within most Congresses such that that the nearest neighbor method cannot identify suitable matches for the small number of LGB legislators. Although we pool all congressional sessions together, we do require that legislators are only matched within the same congressional session. Further, given the small number of Republican LGB legislators, we only conduct this analysis for Democratic legislators. We estimated the SATT with varying numbers of minimum matches from 3 to 7, because while the optimal number is not known Abadie et al. (2004) suggest that four matches is reasonable.[3]

This implementation of nearest neighbor matching imputes the potential outcomes for each treated/control observation by using the average of the N nearest neighbors based on the specified covariates. The estimated treatment effect is then the difference between the observed and imputed outcomes across all observations (Abadie et al. 2004). This is different than matching as pre-processing as done in CEM, which reduces the dataset to the best matched observations and then we estimate the treatment effect using OLS to see if there is a difference in the outcomes between the matched treated/control observations. The results in Table D1, are substantively similar to those reported in the

[2] We used the teffects package in STATA, v18.5.
[3] Abadie, A., Drukker, D., Herr, J. L., & Imbens, G. W. (2004). Implementing Matching Estimators for Average Treatment Effects in Stata. The Stata Journal, 4(3), 290-311. https://doi.org/10.1177/1536867X0400400307

main appendix. Furthermore, the estimated effect of LGB status on HRC score is nearly identical to the estimates in the paper, which given the significant difference in both the matching methods and the estimation strategy suggests that our results are not highly sensitive to the exact matching and estimation strategies.

Table D1: Effect of LGB Status on HRC Scores, Using Nearest Neighbor Matching

	(1) Number of Neighbors	(2) Number of Matches	(3) Number of Matches
VARIABLES	3	5	7
Effect of LGB Legislator	4.507***	4.351***	4.353***
	(1.236)	(0.941)	(0.849)
Observations	1,683	1,683	1,683

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1