

Incivility and Electoral Outcomes

Jack Banks, Class of 2027

Now more than ever, political campaigns are riddled with incivility and personal attacks. But does this incivility benefit candidates at the voting booth? We put this question to the test through an analysis of over 300 gubernatorial campaigns spanning from 2010-2023. We had two main empirical challenges: the first was to rate the incivility of each candidate and the second was to control for all other variables that could influence electoral performance. Below is a summary of our methodology and key results.

We began by collecting debate transcripts from cspan.org and applying Google Gemini 1.5 Flash to rate the polarity of each candidate's debate performance based on UNITE's Dignity Index, a validated score of polarity. However, due to the tediousness of collecting debate transcripts, this data only spanned back to 2018. Additionally, we felt that debate transcripts were an incomplete representation of a candidate's campaign.

To address both shortcomings, we asked three AI models (GPT-4o, Claude Sonnet, DeepSeek V3) to rate the incivility and frequency of character attacks employed across a candidate's entire campaign. This differed from our first approach because we did not provide the LLM with any material, instead requesting it to assign scores based on all available knowledge. The scores from these three models had strong correlations with the initial debate ratings, thus supporting the validity of the LLM measures. Seeking further validation of the three models, we used Google's Perspective API to rate the debate transcripts for toxicity and again saw reasonably high correlations between measures.

We then controlled for variables beyond rhetoric using a forecasting model from Love et al. that used ex ante data to forecast gubernatorial elections. Because our goals differ slightly from theirs, we also included ex post variables like third party vote share (if greater than 5%) and the changes in state legislative seats for the incumbent's party. We excluded approval ratings from the forecasting model on the basis that rhetoric could influence approval. Two additional controls—ideological extremism and populism—were also rated by the three LLMs and validated against external sources.

The key independent variables were the incumbents mean LLM incivility score minus the challenger's and the incumbents mean LLM character attack score minus the challenger's. The dependent variable was incumbent vote percentage. Linear regressions showed that for every one-point increase in a candidate's civility (or decrease in character attacks) there is a 2.5-3 percentage point gain in vote share. Civility effects were significant at the 10% level; character attack effects at the 1% level.

To strengthen a causal interpretation, we also examined “underdog effects”, whereby candidates who are behind in the polls use more character attacks to gain ground, and that the reason for the positive association seen above is that uncivil candidates were already losing. However, using 3,782 state polls we found that even when candidates were equally behind in the polls, those who maintained more civil rhetoric consistently outperformed their less civil counterparts. Thus, this suggests that even when controlling for polling, incivility does not produce gains.

These findings suggest that the assumption that “going negative” is electorally advantageous may be misguided. While incivility might energize certain partisan bases, the broader electorate appears to penalize it. The 2.5-3 percentage point swing associated with civility is substantial, especially in close races where such margins often determine the outcome. This has implications for not only campaign strategy but also the health of democratic discourse: candidates have an electoral incentive to run more respectful campaigns. These results offer a counter-narrative, suggesting that civil discourse may be both normatively desirable and strategically advantageous.

Faculty mentor: Daniel Stone

Funded by the Breckinridge Fund

Sources:

Love, G. J., Carlin, R. E., and Singer, M. M. (2025). Lassoing the governor's mansion: A machine-learning approach to forecasting gubernatorial elections. *PS: Political Science & Politics*, 58(2):226–233.

Sayed, H. (2024). Persuasion with partisanship: The informational content of policymaking with application to us governors.