

Blameocracy: Causal Attribution in Political Communication*

Francesco Bilotta

Alberto Binetti

Giacomo Manferdini

April 9, 2025

Abstract

We propose a supervised method to detect causal attribution in political texts, distinguishing between expressions of merit and blame. Analyzing four million tweets shared by U.S. Congress members from 2012 to 2023, we document a pronounced shift toward causal attribution following the 2016 presidential election. The shift reflects changes in rhetorical strategy rather than compositional variation in the actors or topics of the political debate. Within causal communication, a trade-off emerges between positive and negative tone, with power status as the key determinant: government emphasizes merit, while opposition casts blame. This pattern distinguishes causal from purely affective communication. Additionally, we find that blame is associated with lower trust in politicians, perceived government effectiveness, and spreads more virally than merit.

Keywords: Narratives, Text-as-Data, Negative Campaigning, Affective Polarization, Voting

*Francesco Bilotta: Bocconi University, francesco.bilotta2@phd.unibocconi.it; Alberto Binetti: Princeton University, abinetti@princeton.edu; Giacomo Manferdini: Bocconi University; giacomo.manferdini@phd.unibocconi.it. We are grateful to Enrico D. Turri for his input in earlier stages of this project. We thank Luca Braghieri, Sarah Eichmeyer, Nicola Gennaioli, Matteo Grigoletto, Dirk Hovy, Rafael Jiménez-Durán, Massimo Morelli, Carlo Schwarz, and Ekaterina Zhuravskaya for their helpful comments and suggestions, as well as participants to Munich Workshop on Beliefs Narratives and Memory and 2nd Bocconi-CCA-Cornell Political Economy Workshop.

A well-established finding in the psychology of reasoning is that humans are innately drawn to causal explanations (Chater and Loewenstein, 2016; Lombrozo, 2012; Sloman and Lagnado, 2015). Correspondingly, recent economic evidence shows that when information is woven into “narratives” it is more effective in shaping people’s beliefs (Alesina et al., 2023) and decisions (Hüning et al., 2022). At the same time, compared to other types of statements – such as moral imperatives or affective messages – causal explanations embody an additional constraint, as they need to be credible, at least to some extent (Ambuehl and Thyssen, 2024; Barron and Fries, 2024a).¹

Arguably, the tension between persuasiveness and credibility is especially acute in political communication, where candidates must convince voters they are better suited than their opponents to govern. In this context, it is well known that politicians appeal to voters’ sentiment, leveraging affective polarization as a means to reinforce party identification (Ansolabehere and Iyengar, 1995; Iyengar et al., 2012; Sood and Iyengar, 2016, e.g.). At the same time, politicians may improve on purely affective communication by sharing causal explanations that *attribute* favorable outcomes to their own actions – taking merit – while unfavorable ones to their opponents’ – shifting blame.

Hence, it is a natural question to quantify the supply of causal attribution in political rhetoric, and to understand whether causality imposes distinctive rules and constraints on communication, compared to affection.

To address these issues, we propose a supervised method to detect causal attribution in political text, distinguishing between expressions of merit and blame. We apply our method to analyze a large corpus of tweets shared by U.S. Congress members between 2012 and 2023. Our main finding is a pronounced shift toward causal attribution following the 2016 presidential election. We show that this shift is the outcome of an active change in rhetorical strategy, rather than the mechanical consequence of compositional variation in actors or topics of the political debate. Moreover, within causal communication, a clear trade-off emerges between positive and negative tone, with power status emerging as the key determinant – the governing party claims merit while the opposition casts blame. Notably, this pattern is distinctive of

¹Beyond persuasiveness, other factors may contribute to the benefits of spreading narrated information. Compared to facts, opinions are more likely to be actively sought out (Bursztyn et al., 2023) while stories are more likely to be remembered (Graeber et al., 2024). Explanations increase the likelihood that public choices are imitated (Graeber et al., 2024) while rationales may act as a “social cover” that fosters the public expression and tolerance of stigmatized positions (Bursztyn et al., 2023). Finally, narratives increase the subjective valuation of items (Morag and Loewenstein, 2024).

causal attribution – it does not emerge for purely affective messages.

The primary challenge we face is measuring causal attribution. Standard computational linguistics tools often struggle to detect causality – let alone its interplay with tone – because causal cues are typically conveyed implicitly rather than through a fixed set of semantic or syntactic markers. Consequently, we employ a supervised learning approach based on bidimensional classification.

Our method is grounded in a definition of causal attribution inspired by the Neyman-Rubin causal model (Neyman, 1923; Rubin, 1974). Specifically, we define a tweet as *causal* if it attributes a potential outcome to the (hypothetical) intervention of a political agent. Additionally, we classify the tweet’s *tone* as negative, positive, or neutral, based on the attitude expressed toward its subject matter. Within causal tweets, we define those with positive tone as *merit* and those with negative tone as *blame* – we refer to this feature as the tweet’s rhetorical style.

Based on this definition, we tag a training dataset of approximately 4,000 tweets by rhetorical style. Despite the complexity of detecting causal attributions, inter-annotator agreement remains high – with a Fleiss’ Kappa of 0.64. Then, we train a multi-class classifier to predict whether a tweet expresses merit, blame, or none, by fine-tuning a RoBERTa-large model pre-trained on 154 million tweets (Loureiro et al., 2022). We then deem a tweet causal, ex-post, if its predicted style is either merit or blame.² Our classifier achieves robust performance, meeting or exceeding standard benchmarks in the literature – with an accuracy of 0.83, an F1-score of 0.84, and a Matthews Correlation Coefficient of 0.73.

Our measure of merit and blame passes a set of natural validation checks. For internal validity, we first focus on the referents of merit and blame texts. Intuitively, blame texts are predominantly other- or out-group referential, whereas merit texts are chiefly self- or in-group referential. This pattern holds both syntactically, as shown by pronoun counts, and semantically, when targets are identified via the Political DEBATE language models (Burnham et al., 2024). Second, analyzing temporal tenses shows that blame tweets are more retrospective while merit tweets are more prospective – a pattern consistent with the intuitive idea of “borrowing” future good outcomes, while distancing from past negative ones. Both results are further corroborated by a qualitative bigram analysis. For external validity, we compare our measure with the metrics of “credit-claiming” and “policy-attack” statements from America’s Political Pulse (Westwood et al., 2024), finding a strong correlation at the individual politician level.

Finally, to further support the novelty of our blame-merit dimension, we show

²This choice is justified by the fact that, in our training dataset, virtually no causal tweets exhibit a neutral tone. To validate this approach, we also train an independent classifier for causality alone, showing high correlation with our synthetic label. Compared to having two separate classifiers, this approach allows us to minimize classification noise.

that it does not reduce to a linear combination of established measures for various linguistic features of the political discourse, such as sentiment (Hutto and Gilbert, 2014), emotionality (Gennaro and Ash, 2022), and moral vocabulary (Enke, 2020).

Our conceptual innovation is instrumental in addressing our empirical questions. We begin by examining politicians' reliance on causal communication. Our main finding is a striking upward trend in the share of causal tweets during our period of analysis. The fraction of causal tweets more than doubles, for both parties – moving from an average of 20% during Obama's presidency (2012-2016) to an average of 40% during Biden's (2020-2024). Notably, most of this increase occurs immediately after the 2016 election, driven primarily by Democrats "turning causal," while Republicans follow with some lag during Trump's term. Importantly, none of the other textual dimensions considered in our analysis exhibits a comparable pattern, as shown by structural break and seemingly unrelated regression analysis.

Given this striking evidence, one might wonder if the shift to causality is the byproduct of exogenous dynamics undergone by the political debate – such as compositional changes in the pool of politicians or in the pool of topics they discuss. Controlling for politician and topic fixed effects, we rule out both hypotheses, demonstrating that the shift takes place within-politician and within-topic. Hence, we interpret this increase in causality as reflecting a change in rhetorical strategy, rather than a mechanical phenomenon. Moreover, the increase in causal language is stronger among Democrats and in policy-related (rather than sociocultural) issues, while no demographic feature is significantly correlated with the shift.

If adopting causal language is a strategic choice, it is interesting to examine the distinctive features of this kind of rhetorical competition and how it differs from one based on sentiment only. Our analysis reveals that after politicians shift to causal communication, a clear substitution pattern emerges between merit and blame. To understand what underlies it, we study the determinants of a politician's emphasis on merit or blame, within their causal tweets.

We show that, while fixed politician-level characteristics (such as demographics) do not strongly influence this margin, power status – which is inherently dynamic – is the key predictor. In particular, the governing party (i.e., the party of the President) always claims more merit while the opposition shifts more blame – an effect which increases through time. Importantly, power status does not have a comparable influence on sentiment within non-causal tweets. This finding suggests that the choice between merit and blame is a strategic trade-off rather than a fixed politician trait, and that it is distinctive of the causal domain. In particular, the trade-off seems to reflect a minimal accountability principle, whereby those in power cannot offload all blame, while those in opposition cannot claim all merit.³ Interestingly, members of

³This pattern aligns with the key predictions of the partial identification approach to narratives

different parties adjust to these constraints along different margins, as Republicans do relatively more blame when in opposition, while Democrats do more blame when in power.

Overall, our findings show that causal attribution is increasingly important in politics and that it does respect intuitive and distinctive regularities. As a bridge towards further research, we present suggestive evidence in three directions. First, we analyze all the newsletters sent by Congress members to their constituents. We find a similar pattern – an increase in the usage of causal sentences – around the same period, suggesting that our findings are meaningful also outside of Twitter. Second, we find that politicians’ rhetorical style appears to match voters’ attitudes. In particular, we find the amount of causal attribution to be positively associated with the level of affective polarization, and the amount of blame to predict lower trust in politicians and lower perceived government efficiency. Finally, in a case-study focusing on “presidential tweets” for Obama, Trump, and Biden, we show that blame tweets are significantly more viral than merit ones.

The rest of this paper is organized as follows. We conclude the Introduction discussing some related literature. We introduce our definition and classification procedure in Section 1, while in Section 2 we present our validation checks. Our empirical findings are shown next: Section 3 explores the shift to causality, whereas Section 4 the merit-blame tradeoff. We conclude with our discussion in Section 5.

Related Literature

Our paper contributes to the literature on text-as-data in political economy. Conceptually, we propose a novel definition of causal attribution, merit-taking, and blame-shifting. In this sense, we add to previous research that measures other key linguistic elements of political language, such as emotionality (Gennaro and Ash, 2022), moral terminology (Enke, 2020), and complexity (Di Tella et al., 2023). A recent paper closely related to ours is Algan et al. (2025). Analyzing both voters’ and politicians’ policy-related tweets from 2013 to 2024, the authors document a marked rise in negative emotions – especially anger – starting in 2016. We view their findings as complementary to ours in two respects. First, the timing of this emotional shift aligns closely with the surge in causal attributions we uncover, reinforcing the interpretation that a structural change in political rhetoric took place post-2016. Second, the prominence of anger, relative to other emotions, provides a psychological rationale for the rise in causal communication. Among negative emotions, anger is uniquely directional: it implies a belief that someone is responsible for a perceived harm, and thus naturally calls for a causal explanation (Lazarus, 1991). In this sense, the growing presence

proposed in Bilotta and Manferdini (2024).

of anger may help account for the increasing use of causal language in political discourse.

Methodologically, we contribute to the literature on text-as-data methods in economics (Ash and Hansen, 2023; Gentzkow et al., 2019, for a review) by introducing a supervised learning approach for causality detection. This approach departs from the existing methods in the literature, which are mainly based on unsupervised dictionary- or part-of-speech-based algorithms. Detecting causal attributions is notoriously difficult as it does not always hinge on the usage of specific words.⁴ The unsupervised method most closely aligned with our purpose is Ash et al. (2024), as it allows to detect explicit agent-verb-patient representation. Nonetheless, its application poses two non-trivial issues. First, the method should be tuned to detect, among all representations, only causal attributions. Second, even after solving this issue, implicit causal attributions would still elude the classification. In contrast, our supervised approach leverages the broader contextual understanding of language models to recognize both explicit and implicit causal arguments more accurately. In this sense, our strategy resonates with the spirit of methods used to detect the slant or sentiment in newspapers (Braghieri et al., 2024; Shapiro et al., 2022) or the politicians’ appeal to populist rhetoric (Bellodi et al., 2023). Overall, the advantages of this approach – especially using BERT models – to detect concepts in texts are detailed in Ash and Hansen (2023, Section 3.2.2.4).

We also contribute to the literature on narratives (Barron and Fries, 2024b, for a review). Much of the existing empirical research relies either on surveys (Andre et al., 2021, e.g.) or experiments (Ambuehl and Thyssen, 2024; Barron and Fries, 2024a; Kendall and Charles, 2023, e.g.). We innovate on this margin, adding to a small number of studies that measure narratives using observational data (Gehring and Grigoletto, 2023; Goetzmann et al., 2022; Macaulay and Song, 2023). We differ from these works in several ways. First, our approach is domain-agnostic, while most works focus on specific topics. Second, we capture both the extent and the direction of causal attribution without imposing any structure on how causality is expressed or which agents are involved – relying instead on the text’s inherent semantic content. These advantages are made possible by the complementarity between our definition, which is inspired by theoretical work on narratives (Aina, 2021; Eliaz and Spiegler, 2020; Ispano, 2025; Jain, 2023; Schwartzstein and Sunderam, 2021), and our methodology. Our definition is particularly close in spirit to the notion of optimal plausible narratives developed by Bilotta and Manfredini (2024), and also captures scapegoating dynamics discussed by Eliaz et al. (2023). Finally, while we do not focus specifi-

⁴For instance, “causality” is one of the worst-performing labels in LIWC-22 (Pennebaker et al., 2022) – a gold standard for dictionary methods; as Figure A1 shows, its Cronbachs α ranks among the lowest.

cally on the moral domain, the competition between causal attributions and affective messages in our framework resonates well with that between moral imperatives and narratives in [Bénabou et al. \(2018\)](#).

A vast literature in economics and (especially) political science studies the determinants and evolution of affective polarization ([Boxell et al., 2024](#); [Iyengar et al., 2012](#)); see ([Iyengar et al., 2019](#)) for a review. While much of this work focuses on mass attitudes, some studies examine affective polarization among political elites and its downstream consequences for voters ([Broockman and Butler, 2017](#); [Diermeier and Li, 2019](#)). Moreover, recent contributions argue that affective polarization has come to play a larger role than ideological polarization ([Baldassarri and Page, 2021](#); [Druckman and Levy, 2022](#)). We contribute to this debate by introducing the notion of a “claimed” causality margin, which may help explain why politicians continue to reference the opposing party even when such messages offer no clear electoral benefit ([Costa, 2021](#)). While these references may appear expressive in aggregate, distinguishing causal claims may reveal a more strategic rationale behind them.

Relatedly, we contribute to the literature on negative campaigning, which emphasizes candidates’ decision between running on their own strengths (a “positive” appeal) or to concentrate on exposing the weaknesses of their opponents’ (a “negative” appeal) ([Lau and Rovner, 2009](#), for a review). Our findings quantitatively show that causal attribution plays an increasingly important role in campaigning, and that this constrains the decision between positivity and negativity in a meaningful fashion. Moreover, our suggestive evidence on the virality of blame aligns well with studies arguing for a perceptual motivation for negativity ([Lau, 1985](#); [West, 2018](#), e.g.), compatibly with psychological evidence of negativity bias in information processing ([Baumeister et al., 2001](#); [Rozin and Royzman, 2001](#), e.g.).

Finally, we relate to the study of retrospective voting in political science ([Healy and Malhotra, 2013](#), for a review). Within this literature, a classic strand shows how voters attribute blame to political actors for events often beyond their control ([Converse, 1964](#); [Malhotra and Kuo, 2008](#), e.g.), thus incentivizing politicians to make suboptimal decisions to avoid it ([Weaver, 1986](#)). Importantly, this attribution process is often group-serving and hence prone to partisan bias ([Taylor and Doria, 1981](#), e.g.). We complement this perspective empirically by showing that politicians often share (potentially false) narratives in order to shape the attribution process. In this sense, we measure the “blame-game” pictured by [Stone \(1989\)](#) and [Hood \(2010\)](#).

1 Measuring Causal Attribution

In this section, we detail our approach for detecting causal attribution in political communication. First, we describe the corpus of political tweets collected in our dataset. Next, we define causality, merit, and blame in our context. Finally, we outline our classification procedure.

Data Our primary dataset comprises approximately 4 million tweets posted by Democratic and Republican U.S. Congress members between 2012 and 2023. We combine data from the CongressTweets project ([CongressTweets, 2023](#)) and from [Belodi et al. \(2023\)](#). To the best of our knowledge, the dataset includes all tweets posted by House representatives from January 3, 2013, through July 11, 2023; all tweets by House representatives running for re-election in 2012 (covering January 1 through December 31, 2012); and all tweets posted by Senators from June 21, 2017, through July 11, 2023. In total, our sample contains 4,198,455 tweets from 900 Congress members across 1,789 Twitter accounts.⁵ We also enrich the dataset with demographic and ideological information from ProPublica, VoteView, and Wikipedia. Table A1 presents descriptive statistics at both the tweet and politician levels. In our sample there is a small majority of Republicans politicians, but the majority of tweets are instead posted by Democrats who, as found also by [Fujiwara et al. \(2024\)](#), are more active on Twitter. Most politicians have a bachelor’s degree or higher, the average age of the user is slightly below 60, and only about a quarter of them are female.

Definition of Causality, Merit and Blame Our supervised learning approach employs a bidimensional classification, assigning each tweet labels for *causality* and *tone*.

The first dimension, causality, is binary. A tweet is deemed *causal* if it attributes a potential outcome to the power status of a political agent. Political agents include politicians, political institutions, politically aligned organizations, and other individuals known to influence the political discourse. They exclude natural events – such as pandemics or disasters – as well as neutral agents – such as scientific teams. Power status refers to an agent’s capacity for significant political action (for instance, stimulating the implementation of policies or reforms). Power status can be factual – when the statement implies the agent is currently taking or has previously taken action – or hypothetical – when it implies that the agent could take action if in power. We impose no restrictions on the potential outcome dimension – this may encompass any feature of the state of affairs experienced by the polity, ranging from economic outcomes, such as GDP growth, to social outcomes, such as civil liberties. Finally, we

⁵We exclude retweets as they do not represent original content, while we instead keep quoted tweets as they also contain content generated by the user.

stress that, in our approach, the attribution of causality does not need explicit causal connectors (such as “because” or “since”).⁶

The second dimension, tone, categorizes tweets into three levels: negative (-1), neutral (0), and positive (1). Tone reflects the emotional or attitudinal stance of the tweet toward its subject matter. A tweet is considered positive or negative if its language explicitly or implicitly conveys favorable or unfavorable sentiments, respectively. Neutral tweets are those that objectively describe or discuss their subject matter without evident emotional or evaluative language. Our manual annotation captures subtle nuances that traditional dictionary methods (e.g. [Hutto and Gilbert \(2014\)](#)) might overlook.⁷

Considering the product of causality and tone, we create a synthetic measure termed *rhetorical style*, which assumes three distinct values: merit (1), blame (-1), or none (0). Tweets identified as merit (blame) are those that positively (negatively) attribute outcomes to political agents’ power status. All remaining tweets are labeled as none, indicating either non-causal statements or causal statements with a neutral tone. To get a sense of our definition, [Table A4](#) provides labeled examples.

Labeling Procedure Leveraging these definitions, we manually annotate a training dataset consisting of 3,958 tweets. To ensure balanced representation across classes, given that merit and blame tweets may be relatively rare, we implement a two-step oversampling strategy. First, we instruct ChatGPT to generate 50 example tweets for Democrats and Republicans that reflect our definition of merit and blame. Then, for each tweet in the main dataset, we compute cosine similarities with these examples using SBERT mini embeddings ([Reimers and Gurevych, 2019](#)). Half of the tweets selected for annotation are those with the highest similarity to the generated examples (balanced equally between merit and blame), while the remaining half are randomly sampled. Additionally, 530 tweets are randomly chosen for joint annotation by three coders to assess inter-annotator agreement. The final labels for causality and tone are determined by majority voting, with ties broken randomly. Our annotation protocol achieves an average pairwise correlation of 0.73 and a Fleiss’ Kappa of 0.64 – generally interpreted as substantial agreement. Overall, the coders agree on the rhetorical style in 67% of tweets – six times the rate expected by chance.

⁶Overall, the major distinction causality operates is between mere moral or evaluative judgments – such as “Politician X is corrupt” – which do not satisfy the criterion for causal attribution – and statements that do imply a link from actions to consequences – such as “Politician X’s corruption undermines democracy”.

⁷For instance, a sentence like “Our policies avoided a tragedy” would be assigned negative sentiment, but has positive tone according to our annotation protocol.

Fine-Tuning Drawing on our annotated dataset, we proceed to fine-tune a RoBERTa large model pre-trained on 154 million tweets (Loureiro et al., 2022).⁸ In particular, we split the labeled corpus into an 80 percent training set (3,166 tweets) and a 20 percent validation set (792 tweets).⁹

Classifier Output and Performance Our classifier outputs, for each tweet, a distribution over three categorical classes – Merit, Blame, and None. As standard, the tweet is assigned to the class that has the highest probability, that is, the mode of the output distribution. This choice could be problematic in instances where the classifier is “unsure”, i.e., it assigns high probability to more than one category. Nonetheless, these cases are extremely rare for our classifier, validating our approach.¹⁰ In our dataset, merit tweets are about 20 percent, whereas blame tweets are about 16 percent. Taken together, over one-third of tweets fall within our dimension of interest. On the validation set, our classifier achieves an accuracy of 0.83, an F1-score of 0.84, and a Matthews Correlation Coefficient of 0.73, demonstrating its robust performances.

Mapping Output to Variables Our analysis is based on two derived variables. The first, *Causality*, is a binary indicator that denotes whether a tweet is assigned to one label between Merit or Blame. Our interpretation is motivated by the extremely low incidence of causal tweets with neutral tone in our training data (0.4%). To support it, we train a dedicated classifier for causality, based on the corresponding label in our training data¹¹. Then, we compute the correlation between its binary output and Causality. As shown in Figure A4, the correlation is extremely high ($\rho = 0.9$). Moreover, less than 6% of the tweets labeled as causal by the dedicated classifier are not classified as either merit or blame by the main classifier, supporting our initial intuition.¹² The second variable of interest, *Blame*, records whether a tweet is Blame (1) or Merit (0), conditional on being either one of the two – that is, it takes missing value (\emptyset) when Causality is 0. The purpose of these variables is complementary, as they

⁸Fine-tuning involves adapting a pre-trained language model to a specialized downstream task – such as a classification task – by adjusting its parameters on a smaller and specific dataset. This procedure is common in employing language models, and often leads to significant performance gains relative to training a model from scratch (Devlin, 2018; Liu, 2019). In our context, fine-tuning allows the classifier to better capture the nuances of political discourse, such as causal attributions and the tone of these, thereby improving its predictive accuracy.

⁹We fine-tune the model for 10 epochs, selecting the epoch with the highest F1 score to mitigate the risk of overfitting. As an additional safeguard, we also fine-tune both a version of the models with half of its layers blocked and the widely known BERTweet model (Nguyen et al., 2020). In both cases, aside from marginally lower performance, the classification output is analogous to ours, with Matthews Correlation Coefficients of 0.93 and 0.85, respectively.

¹⁰In Figure A3, we plot the density of the output probability that a tweet belongs to a certain class, conditional on being assigned to it. All these densities are strongly left-skewed.

¹¹The training procedure is the same used for the classifier for merit and blame. The causality classifier has an accuracy of 0.83, comparable to the merit-blame one.

¹²Finally, results based on *Causality* hold unchanged if we use the output of the dedicated classifier.

capture, respectively an extensive and an intensive margin in our dimension of interest. On one hand, causality assesses how prevalent causal attribution is in politicians’ tweets. On the other hand, we use our Blame measure to understand the trade-off between these two dimensions, within causal tweets. On occasion, we also report results looking at blame and merit across all tweets, not just causal ones. In these cases, blame and merit are two separate indicators taking binary values. Throughout the text, we refer to these variables as Unconditional Blame and Unconditional Merit.

2 Validation

This section reports the validation exercises performed to support our measure. These are organized in three parts. First, we provide a qualitative description of the text sub-corpora classified as merit and blame, respectively. Second, we benchmark our measure against conceptually related metrics from America’s Political Pulse (Westwood et al., 2024). Finally, we show that merit and blame identify a dimension of political language which cannot be reduced to a combination of the ones measured in the literature.

Qualitative Features of Merit and Blame Text We show how blame and merit texts display both syntactical and semantic features in line with an intuitive understanding of these concepts.

First, we focus on their target. Syntactically, we measure whether text is mostly self- or other-referential, employing a simple vocabulary-based approach. In particular, we consider the difference in count between second- and third-person pronouns versus first-person pronouns, normalized by the number of words in each tweet. Figure A7a shows that, as expected, blame-oriented tweets tend to be framed around others, whereas merit-oriented tweets feature a stronger focus on the self. Semantically, we investigate the political affiliation of the targets of causal attribution. Leveraging the Political DEBATE language models (Burnham et al., 2024), we classify tweets based on whether they are about the Democrats or Republicans. Restricting our focus to the 40 percent of tweets where the party can be recognized by the language model, we show in Figure A7b that blame tweets are substantially more likely to be about the opposing party compared to merit tweets.

Second, we turn to the temporal dimension of causal attributions. In our classification, causal arguments can be retrospective – attributing responsibility for outcomes that have already occurred – or prospective – attributing potential consequences to present actions. To capture this distinction, we measure how frequently past and future tenses appear in each tweet, normalize these counts by the total number of words, and subtract the latter from the former. We illustrate in Figure A7c that blame

tweets typically adopt a more retrospective stance, whereas merit tweets more frequently assume a forward-looking perspective. This retrospective nature of blame is in line with the psychology literature (Malle et al., 2014).

Finally, we follow Gentzkow and Shapiro (2010) to identify the bigrams that most distinctively characterize merit and blame tweets for both parties. As shown in Figures A5 and A6, the bigrams align with our intuitive expectations and findings from the previous two points.¹³ To conclude, in Tables A5 and A6, we present a selection of tweets categorized as conveying merit or blame for Democrats and Republicans, illustrating concretely how the classifier works.

External Validation: America’s Political Pulse We compare our indicators with data from America’s Political Pulse (Westwood et al., 2024). This project, starting from the last quarter of 2022, tracks a range of rhetorical dimensions across multiple forms of political communication, including Twitter posts. We focus on their metrics for “credit-claiming” and “policy-attack” statements, which are the closest proxy to merit and blame.¹⁴ For each politician in the America’s Political Pulse dataset, we construct two measures analogous to our Causality and Blame. In Figure A2, we plot the binscatters of these analogous measures against our Causality and Blame-Merit measures during the same period. In both cases, we find a strong positive association, with correlation coefficients of 0.51 and 0.72, respectively.


Merit and Blame as a Distinct Semantical Dimension Finally, we investigate how our measures relate to well-established indices for other linguistic dimensions of relevance in the political discourse. Specifically, we compare Causality and Blame to sentiment (Hutto and Gilbert, 2014), emotionality of language (Gennaro and Ash, 2022), and the prevalence of moral terminology (Enke, 2020).¹⁵ More precisely, we

¹³Among Democrats, for example, “act will” is diagnostic of merit tweets, highlighting how a politician’s action is credited for future outcomes, while “trump administration” is diagnostic of blame tweets, reflecting how Democrats frequently faulted the Trump administration. Among Republicans, “act will” again surfaces as indicative of merit tweets, whereas “southern border” emerges as diagnostic of blame tweets, mirroring Republicans’ tendency to blame Democrats for immigration issues.

¹⁴In their codebook, Westwood et al. (2024) define “credit claiming” a communication about *Creating or passing legislation; Securing government spending, grants, or funding; Emphasizing personal or party accomplishments in office* and “policy attack” a communication about *Objecting to or raising concerns about a specific policy, law, or court ruling; Using fact-based arguments, even if critical or negative; Avoiding emotional appeals, inflammatory language, claims of extremism, or personal attacks on individuals involved with the policy including accusing them of lying or withholding information..* More information about this source is available [here](#).

¹⁵To measure the emotionality of language and the prevalence of moral terminology we follow the same procedures outlined in the respective papers. Concerning emotionality, we produce embeddings of emotional and reasoning words, and for each tweet we compute the ratio of the cosine similarity with the emotional embedding to the cosine similarity with the reasoning embedding. For the prevalence of moral terminology, we compute, for each moral value, the average between the average frequency of vice and virtue words, then, we sum over all moral values and normalize by the total

Figure 1: Correlation with Existing Text Measures



Notes: Panel (a) presents the estimates from the regression at the tweet level of Causality over the listed text features. Panel (b) presents the estimates from the regression at the tweet level of Blame over the listed text features. In all panels regressors are standardized. In all Panel bars denote 95 percent confidence interval with standard errors clustered at the politician level.

estimate the tweet-level regressions

$$y_{i,p} = \alpha + \beta_1 S_{i,p} + \beta_2 E_{i,p} + \beta_3 M_{i,p} + \mu_p + \varepsilon_{i,p}, \quad (1)$$

where $y_i \in \{\text{Causality}_{i,p}, \text{Blame}_{i,p}\}$ is either the Causality or the Blame value for tweet i posted by politician p ; $S_{i,p}$ is the sentiment score of tweet i ; $E_{i,p}$ is the emotionality of the language of tweet i ; $M_{i,p}$ is the moral value score of tweet i ; and μ_p are politician fixed effects. Errors are clustered at the politician level. All regressors are standardized to make estimates comparable.

Figure 1 shows that our measure of causality is largely orthogonal to these linguistic features, with all coefficients below 0.1 in absolute value, suggesting that it captures a distinct rhetorical dimension. The relatively low level of the R^2 – around 0.13 even including politician FE – supports this view. Reassuringly, the strongest text-level predictor of causality is the emotionality of the language, consistent with the idea that causal texts are more reasoning-oriented than emotion-driven. For the Blame dimension, sentiment is its strongest correlate, reflecting how blame (merit) tweets are more likely to contain negative (positive) language. We view this as a useful sanity check. Interestingly, within causal tweets, merit appeals to reasoning more than blame tweets – possibly hinting at a higher burden of proof required to claim credit compared to shifting blame.

number of non-stop words.

3 The Shift to Causality


Having established our conceptual contribution, we now turn to the main empirical finding: the share of causal tweets by U.S. Congress members increases markedly after the 2016 presidential election. This pattern could reflect either a genuine shift in rhetorical strategy or a mechanical consequence of compositional changes – such as turnover in the set of politicians or variation in the topics they discuss. We provide evidence in favor of the former interpretation. Using our fine-grained measure of causality, we show that the increase occurs both *within individual politicians* and *within topics of debate*. These results motivate our subsequent analysis of the determinants of merit and blame attribution.

Time Trends in Causality Figure 2 displays the quarterly evolution of causal tweets, separately by party. Over our sample period, the share of causal tweets more than doubles. During Obama’s presidency (2012–2016), only about 20% of tweets are causal, whereas during Biden’s term (2021–2023) more than 40% are. Notably, nearly two-thirds of this increase occurred immediately after the 2016 election. This jump is initially driven by Democrats “turning causal,” while Republicans catch up, with some lag, over Trump’s term.

To better appreciate the shift, we benchmark the time series for causality against those for the other textual dimensions considered in our analysis. As shown by Figure A8, none of these other dimensions shows a trend comparable to that in causality. To support this qualitative observation, we run tests for unknown-date structural breaks in each time series. As shown in Table A2, the tests (i) identify an extremely large and significant structural break in causality just after the 2016 election and (ii) do not find a comparable break for any other dimension. As a final test, we estimate a Seemingly Unrelated Regressions model, regressing each measure on a time dummy with value one after Trump’s election. Table A3 reports the differences in estimated coefficients, reinforcing the intuition that the shift in causality is much larger than the shift in any other text measure.

Overall, our findings indicate a structural shift in political communication, with the 2016 election serving as a catalyst for the widespread adoption of causal framing, based on the attribution of merit and blame. Although our analysis is not aimed at identifying the exact cause of this shift (see the discussion in Section 5 for some preliminary steps in this sense), we contend that the shift to causality reflects an active, strategic choice by Congress members rather than the passive outcome of mechanical factors. The remainder of this section presents evidence in support of this interpretation.

Figure 2: Time Trend: Causality



Notes: The figure presents the average Causality at the quarter level separately for Democrats and Republicans. Shaded areas denotes 95 percent pointwise confidence intervals.

Ruling Out Compositional Changes The shift in causality may be due, in principle, to compositional changes in the pool of politicians observed in our sample. First, because we lack data on Senators’ tweets during Obama’s presidency, an immediate concern is that systematic differences between House and Senate drive the shift. We rule this out by replicating our analysis using only House members (Figure A9a). A second mechanism at play may be selection: politicians with a causal communication style may become increasingly likely to be elected, over time. We show in Figure A9b that the rise in causality is not affected by restricting to politicians who served continuously throughout the period. Finally, it could be that some group of Congress members is systematically more prone than others to tweet in a causal style. We rule this out showing that none of our demographics systematically predicts causality (Figure A12).

To demonstrate more rigorously that the shift to causality stems from within-politician changes in communication style, we estimate, separately for Democrats and Republicans, the following regression:

$$\text{Causality}_{i,p,q} = \alpha + \sum_{k=2012q2}^{2023q3} \beta^k D_{i,p,q}^k + \mu_p + \varepsilon_{i,p,q}, \quad (2)$$

where $\text{Causality}_{i,p,q}$ is the Causality of tweet i , posted by politician p , in quarter q ; $D_{i,p,q}^k$ is a binary indicator taking value one if the tweet is posted in quarter k ; and μ_p are politician fixed effects. As shown in Figure A10, the pre-2016 election period exhibits a stable baseline, while both Democrats and Republicans significantly increase their use of causal language starting from 2017. Interestingly, in both elections, this

effect is more pronounced for the losing party.

To complete the picture, we ask whether the average within-politician shift to causality is broadly shared across members of Congress or driven by a small set of outliers. To this end, we perform a before-after comparison centered around 2017, separately for each politician who served both before and after the 2016 election. Figure A11 shows that the distribution of these individual-level is clearly skewed to the right, with over 95% of politicians exhibiting a positive change. This confirms that the rise in causal communication reflects a widespread change in rhetorical strategy rather than being driven by a few members.

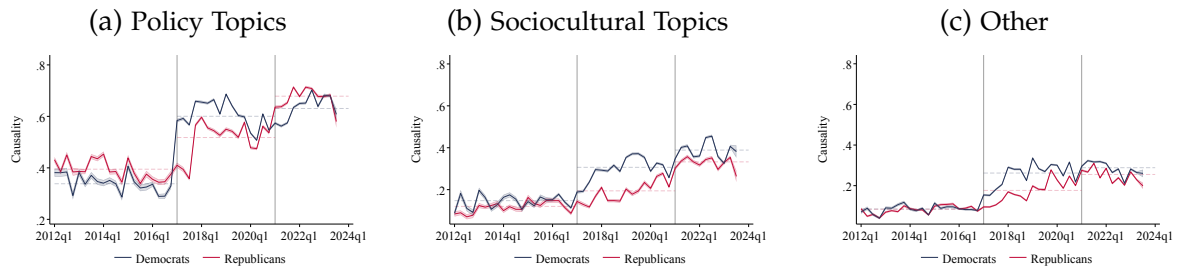
Ruling Out Changes in Topics of Debate While our results suggest that the increase in causal language is primarily a within-politician phenomenon, one might still argue that shifts in the topics of political debate could explain the trend – to the extent that certain topics naturally elicit more causal discussion. To address this possibility, we classify tweets based on the topic of their subject matter using the Political DEBATE language model (Burnham et al., 2024). We group topics into two broad categories: *policy issues* (economy, environment, healthcare, immigration) and *sociocultural issues* (gender, gun control, policing, racial relations). We find that the distribution of topics remains approximately constant over time for both parties (Figure A14), while causal language increases within each topic category (Figure 3a; Figure A13 breaks the results down by topic). Thus, we can rule out changes in subject matter as the driver of our findings.

To refine these insights, we estimate a before-after regression that interacts a post-2016 time indicator with topic dummies, separately by party:

$$\text{Causality}_{i,p,t,q} = \alpha + \sum_{\tau \in \text{Topics}} \beta^{\tau} \left(D_{i,p,t,q}^{2017q1} \times D_{i,p,t,q}^{\tau} \right) + \mu_p + \mu_t + \varepsilon_{i,q}. \quad (3)$$

where μ_p and μ_t denote politician and topic fixed effects, respectively. The results (Figure A15) show that the shift to causal language (i) occurs both within individual politicians and within topic, (ii) is concentrated predominantly on policy issues and (iii) is consistently stronger in Democrats than Republicans, within each topic. Since the pre-2017 distribution of causal language is balanced across parties and topics, points (ii) and (iii) imply that causal tweets are now more prevalent among Democrats and on policy issues overall (Figure A16).

Figure 3: Time Trends: Causality by Type of Topics



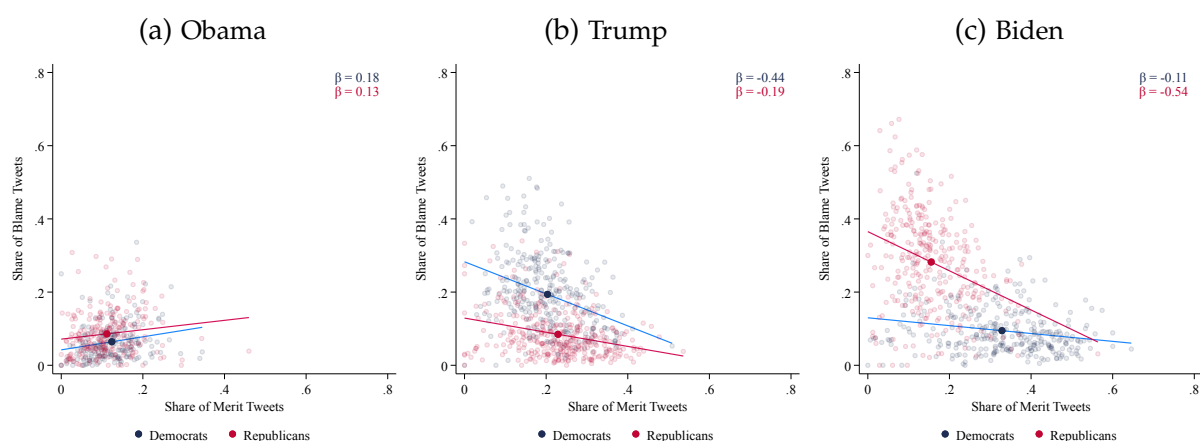
Notes: Panel (a) presents the average Causality at the quarter level separately for Democrats and Republicans restricting to tweets about policy topics. Panel (b) presents the average Causality at the quarter level separately for Democrats and Republicans restricting to tweets about sociocultural topics. Panel (c) presents the average Causality at the quarter level separately for Democrats and Republicans restricting to tweets not classified in any topic. In all panels, shaded areas denotes 95 percent pointwise confidence intervals.

4 Taking Merit or Passing Blame?

The previous Section documented our main finding: a sizable shift toward causal language, observed both within individual politicians and within topics, which seems to reflect the adoption of a rhetorical strategy. We now turn to its internal structure. Our second main result is that, as causal communication grows more prevalent, politicians systematically trade off between merit and blame. A politician’s power status emerges as the strongest predictor of their rhetorical strategy: those in opposition tend to assign blame, whereas those in power emphasize merit. This pattern suggests that, for causal claims to be credible, they must satisfy a minimal plausibility constraint: the relative cost of assigning blame versus claiming merit depends on the politician’s accountability. We support this interpretation through a difference-in-difference analysis showing that the effect of power status on sentiment is significantly stronger for causal than for non-causal tweets.

The Trade-Off between Merit and Blame Figure 4 presents scatterplots of the shares of merit versus blame tweets (that is, computed as a fraction of all tweets) for each Congress member color-coded by party, separately by presidency. A striking contrast emerges when comparing Obama’s presidency (before the shift to causality) with Trump and Biden’s (after the shift to causality). During Obama’s term, no clear relationship between merit and blame is discernible; moreover, the point clouds for both parties substantially overlap. If anything, the positive slope of the regression line suggests that the dominant dimension of heterogeneity during this term is between causal (high merit, high blame) and non-causal (low merit, low blame) politicians. In contrast, following the shift, a distinctive substitution pattern emerges, as indicated by the negative slope in the regression of blame on merit. Figure 4 reveals two key regularities. First, the governing party consistently leans toward merit while the opposition tends toward blame. Second, the trade-off slope is invariably steeper among

Figure 4: Trade-Off



Notes: Panel a presents a scatterplot at the politician level showing the share of tweets classified as merit and blame during the Obama presidency. Panel b presents a scatterplot at the politician level showing the share of tweets classified as merit and blame during the Trump presidency. Panel c presents a scatterplot at the politician level showing the share of tweets classified as merit and blame during the Biden presidency. In all panels the blue (red) solid line is from a linear regression of the share of blame tweets over the share of merit tweets for Democrats (Republicans), with the estimate reported in the top right corner in blue (red). In all panels are excluded politicians who posted less than 10 tweets during the presidency, which are 34, 36, and 5 for the Obama, Trump, and Biden presidency, respectively.

opposition members, suggesting that blame is more “cheaply” employed than merit when in opposition; moreover, this effect becomes more important with time.

Determinants of the Trade-Off Building on the previous evidence, we now ask what drives the trade-off between merit and blame. First, to assess the relative importance of power status as a predictor, we regress the Blame indicator on a set of politician-level characteristics.

The results are shown in Figure A17. We find that demographic characteristics do not significantly predict whether a politician favors blame or merit in their rhetoric. Instead, power status is a strong predictor: members of the opposition consistently exhibit a larger gap between blame and merit tweets, and the strength of the association with power status grows over time.¹⁶ Even within the governing party, members are more likely to assign blame when their party holds a minority in the chamber. Finally, ideological extremism (measured by the absolute Nominat score) is positively associated with blame usage, though this effect declines over time – possibly suggesting a contagion dynamic.

We next ask whether the variation in the blame-merit gap is driven primarily by shifts along the merit margin, the blame margin, or both. We uncover an interesting heterogeneity in the combined effect of power status and ideology. Figure A19 shows the evolution of the share of merit and blame tweets over all tweets in our sample period. Republicans increase their use of blame when in opposition, but show limited

¹⁶Interestingly, this is true even if the overall share of blame tweets is roughly constant over time. To visualize this effect, Figure A18 plots the densities of blame tweets among causal ones.

responsiveness on the merit margin. Conversely, Democrats claim more merit when in power, but do not reduce blame as much when in opposition.

Before-After Comparison for the Effect of Power Status To further support the view that the meritblame trade-off reflects a rhetorical strategy rather than a fixed trait, we estimate two linear models analogous to Equation (2). In the regressions, we use the Blame indicator as the outcome variable and exploit the fact that the first quarters of 2017 and of 2021 mark a shift in power status for each Congress member. Figure A20 presents our results. After Trumps election, Republicans reduce the share of blame tweets among their causal messages by 20%, while Democrats increase theirs by 40%. In the transition to the Biden presidency, the pattern reverses: Republicans’ share of blame rises by 40%, while Democrats’ declines by 30%. In both transitions, the losing party increases blame more than the winning party decreases it.

Figure A21 presents the results stratified by topic. In particular, the effect of power status is significantly stronger on policy issues than on sociocultural ones, reinforcing our accountability argument that causal claims exert greater influence in policy-related debates.

The Effect of Causality on Sentiment As a final point, we show that the meritblame trade-off is a distinct feature of causal communication. To illustrate this, we adopt a difference-in-differences approach comparing the effect of power status on sentiment in causal versus non-causal tweets. First, we construct a synthetic measure of merit and blame by replacing positive and negative tone with positive and negative sentiment, respectively.¹⁷ We then estimate the following model, separately for each presidency:

$$\text{Sentiment}_{i,p,q} = \alpha + \beta \text{Causality}_{i,p,q} + \sum_{X^j \in \text{Demographics}} \beta^j \left(\text{Causality}_{i,p,q} \times X_i^j \right) + \mu_i + \varepsilon_{i,p,q}, \quad (4)$$

where $\text{Sentiment}_{i,p,q}$ represents the sentiment of tweet i posted by politician p in quarter q (with higher values indicating more positive sentiment), $\text{Causality}_{i,p,q}$ is the causality measure of tweet i , μ_p denotes politician fixed effects, and X^j includes demographic variables (e.g., Female, Age, Black, Bachelor, Master or Higher, Opposition, |Nominate Score|, Minority in Chamber). Errors are clustered at the politician level. Our coefficients of interest, the β^j 's, have a difference-in-differences interpretation. For example, $\beta^{\text{Opposition}}$ quantifies the additional effect of being in opposition

¹⁷While our classifier performs well when jointly detecting tone and causality and similarly well when isolating causality its standalone performance on tone is only acceptable. This motivates our choice to proxy tone using sentiment. Although tone captures nuances that sentiment dictionaries may miss, it still correlates meaningfully with sentiment scores.

on negative sentiment, comparing causal tweets to non-causal ones.

Figure A22 shows that the effect of being in opposition on sentiment is significantly more negative for causal tweets than for non-causal ones. Moreover, the estimated effect increases across presidencies, further supporting our interpretation that power status plays a growing role specifically in causal rhetoric. A similar, though noisier, pattern emerges for ideological extremism. Figure A23 plots instead the results separately for causal and non-causal tweets. For causal tweets, our findings largely replicate previous results (albeit with slightly more noise, such as a less precise Nominat score association), further validating our sentiment-based measure of blame. By contrast, non-causal tweets exhibit weaker and less consistent patterns: power status does not robustly predict sentiment, and during Bidens presidency, opposition members even tend to share slightly more positive non-causal tweets.

5 Discussion

We conclude by outlining the boundaries of our analysis and identifying promising directions for future research. In particular, we highlight three themes: (i) establishing the external validity of the observed shift toward causal language; (ii) understanding whether this shift is primarily driven by supply-side strategies or by changes in audience demand; (iii) examining the virality of causal content, and whether messages attributing blame or merit propagate with similar intensity. For each theme, we provide preliminary evidence that supports our arguments.

External Validity: Newsletters A limitation of our analysis is its exclusive focus on one platform, Twitter. Two arguments support this choice. First, Twitter’s focus on the sharing of opinions has made it a standard for studies at the intersection of media economics and political economy (Barberá et al., 2019; Demszky et al., 2019; Halberstam and Knight, 2016, e.g.). Second, by concentrating on communication from politicians to voters, our study naturally limits the range of platforms amenable to analysis – other potential sources include presidential speeches, political ads, manifestos, and political newsletters. As a first step towards external validity, we apply our classifier to the latter, exploiting the dataset collected by Cormack (2025).

This data source presents some conceptual differences from the Twitter data. First, this represents a more tailored communication channel between a politician and their constituency, which means that dialogue is not restricted by platform policies or indirectly influenced by algorithms. Second, for these reasons, this is potentially a more institutional channel of communication.

To apply our classifier, we analyze the newsletter corpus at the sentence-level, so that text-units have comparable length to tweets. Then, for each newsletter piece, we

mimic our tweet classification, first computing the share of merit, blame, and none sentences and then constructing a measure of causality as the sum of the first two shares.

Figure A24 shows that newsletters also exhibit a sharp increase in the proportion of causal sentences written by politicians. Over our sample period, Democrats' causality share increases by 50 percent, while Republicans' by 37 percent. Moreover, once again, much of the shift concentrates around the 2016 election. Compared to the pattern pictured in Figure 2, while the variation is smaller in size, the time series shows the same marked discontinuity. In this regard, it's worth noting that, compared to Twitter, newsletters show an initially higher baseline for causality (30% vs. 20%) – possibly reflecting the more institutional nature of newsletters.

This evidence suggests that the shift to causality may be a domain-general rather than Twitter-specific phenomenon. To strengthen this point, we find that both our measures of causality and blame on Twitter are positively correlated with their analogues in the newsletters data, at the politician-quarter level (Figure A25). This shows that the shift to causality on Twitter and in our newsletter corpus is driven by a common cause – the same politicians adopting a novel rhetorical strategy on both platforms.

Matching between Voters' Attitudes and Politicians' Style It is reasonable to expect that politicians' rhetorical style aligns with voters' attitudes and, conversely, that such attitudes are shaped, to some extent, by the prevailing climate of political debate. While identifying whether demand-side preferences or supply-side choices drive our findings clearly lies beyond the scope of this paper, we provide suggestive evidence that politicians' rhetoric and voters' attitudes are indeed aligned.

To this end, we analyze data from America's Political Pulse (Westwood et al., 2024), this time focusing on the voters' side. In particular, we correlate their survey measures of "People's Beliefs and Attitudes" with the fraction of merit and blame tweets shared by politicians, aggregated at the state-week level. Results are presented in Figure A26. Notably, the prevalence of causal attribution in political communication is associated with stronger ingroup feelings and greater affective polarization – an association primarily driven by the blame margin, in both cases. In addition, both merit and blame correlate in the expected direction with measures of political trust and perceptions of government efficiency, with blame generally showing a stronger association than merit.

If voters' attitudes indeed shape politicians' rhetorical choices, then differences in the voter bases of Democrats and Republicans may help explain the ideological heterogeneity observed in our patterns of interest. In this light, it is useful to relate our findings to established evidence in social psychology. For example, Democrats'

tendency to share more causal tweets (Figure A16) may reflect their voters' stronger consequentialist (rather than deontological) moral orientation (Awad et al., 2020; Graham et al., 2009; Piazza and Sousa, 2014), while Republicans' slant toward blame may mirror a higher negativity bias among their constituents (Hibbing et al., 2014). These intriguing connections merit further exploration, as recent economic research has begun to examine similar questions (Bénabou et al., 2024).

Virality and Causal Attribution At its origins, the literature on narratives has identified virality as a defining feature for the concept (Hirshleifer, 2020; Shiller, 2017). Inspired by this, we ask which text-level feature best predicts tweet popularity. Although Twitter's API restrictions prevent us from collecting engagement measures for our full sample, we conduct a case study using presidential tweets – where retweet data is available. Specifically, we collect tweets from Presidents Obama, Trump, and Biden from various archives¹⁸ and process them with the same pipeline as our main sample.

Figure A27 shows, for each president, how text-level features correlate with tweet popularity. Overall causality does not exhibit a significant association with popularity (Figure A27a). However, this average masks substantial heterogeneity. Across presidential accounts, merit tweets consistently receive fewer retweets than average. In contrast, the effect of blame varies: it is negligible for Obama, but turns positive under Trump and remains positive – though somewhat weaker – under Biden. Whether considering the full sample or restricting to causal tweets only (Figure A27b), the effects of merit and blame on tweet popularity are comparable to, and in some cases larger than, those of other textual attributes such as sentiment or emotionality.

These findings suggest that causal attribution plays a meaningful role in driving tweet virality. In particular, the heterogeneity in the direction of the effect indicates that blame may serve as a tool for attacking opponents and spurring engagement, while merit appears to function more defensively, thus resulting in lower popularity. The dynamics of merit versus blame attribution, and their differential impact on diffusion, warrant further investigation.

¹⁸Tweets from Trump are obtained from the [Trump Twitter Archive](#), from Biden via [Kaggle](#), and from Obama via [Kaggle](#).

Bibliography

- Aina, C. (2021). Tailored stories. Technical report, Mimeo.
- Alesina, A., A. Miano, and S. Stantcheva (2023). Immigration and redistribution. *The Review of Economic Studies* 90(1), 1–39.
- Algan, Y., E. Davoine, T. Renault, and S. Stantcheva (2025). Emotions and policy views. Technical report, Mimeo.
- Ambuehl, S. and H. C. Thyssen (2024). Choosing between causal interpretations: An experimental study. *NHH Dept. of Economics Discussion Paper* (07).
- Andre, P., I. Haaland, C. Roth, and J. Wohlfart (2021). Narratives about the macroeconomy.
- Ansolabehere, S. and S. Iyengar (1995). Going negative: How attack ads shrink and polarize the electorate. (*No Title*).
- Ash, E., G. Gauthier, and P. Widmer (2024). Relatio: Text semantics capture political and economic narratives. *Political Analysis* 32(1), 115–132.
- Ash, E. and S. Hansen (2023). Text algorithms in economics. *Annual Review of Economics* 15(1), 659–688.
- Awad, E., S. Dsouza, A. Shariff, I. Rahwan, and J.-F. Bonnefon (2020). Universals and variations in moral decisions made in 42 countries by 70,000 participants. *Proceedings of the National Academy of Sciences* 117(5), 2332–2337.
- Baldassarri, D. and S. E. Page (2021). The emergence and perils of polarization. *Proceedings of the national academy of sciences* 118(50), e2116863118.
- Barberá, P., A. Casas, J. Nagler, P. J. Egan, R. Bonneau, J. T. Jost, and J. A. Tucker (2019). Who leads? who follows? measuring issue attention and agenda setting by legislators and the mass public using social media data. *American Political Science Review* 113(4), 883–901.
- Barron, K. and T. Fries (2024a). Narrative persuasion. Technical report, WZB Discussion Paper.
- Barron, K. and T. Fries (2024b). Narrative persuasion: A brief introduction encyclopedia of experimental social science.
- Baumeister, R. F., E. Bratslavsky, C. Finkenauer, and K. D. Vohs (2001). Bad is stronger than good. *Review of general psychology* 5(4), 323–370.

- Bellodi, L., M. Morelli, A. Nicolo, and P. Roberti (2023). The shift to commitment politics and populism: Theory and evidence. *BAFFI CAREFIN Centre Research Paper* (204).
- Bénabou, R., A. Falk, and L. Henkel (2024). Ends versus means: Kantians, utilitarians, and moral decisions. Technical report, National Bureau of Economic Research.
- Bénabou, R., A. Falk, and J. Tirole (2018). Narratives, imperatives, and moral reasoning. Technical report, National Bureau of Economic Research.
- Bilotta, F. and G. Manferdini (2024). Coarse memory and plausible narratives. *Available at SSRN 4700043*.
- Boxell, L., M. Gentzkow, and J. M. Shapiro (2024). Cross-country trends in affective polarization. *Review of Economics and Statistics* 106(2), 557–565.
- Braghieri, L., S. Eichmeyer, R. Levy, M. M. Mobius, J. Steinhardt, and R. Zhong (2024). Article level slant and polarization of news consumption on social media. *Available at SSRN 4932600*.
- Broockman, D. E. and D. M. Butler (2017). The causal effects of elite position-taking on voter attitudes: Field experiments with elite communication. *American Journal of Political Science* 61(1), 208–221.
- Burnham, M., K. Kahn, R. Y. Wang, and R. X. Peng (2024). Political debate: Efficient zero-shot and few-shot classifiers for political text. *arXiv preprint arXiv:2409.02078*.
- Bursztyjn, L., G. Egorov, I. Haaland, A. Rao, and C. Roth (2023). Justifying dissent. *The Quarterly Journal of Economics* 138(3), 1403–1451.
- Bursztyjn, L., A. Rao, C. Roth, and D. Yanagizawa-Drott (2023). Opinions as facts. *The Review of Economic Studies* 90(4), 1832–1864.
- Chater, N. and G. Loewenstein (2016). The under-appreciated drive for sense-making. *Journal of Economic Behavior & Organization* 126, 137–154.
- CongressTweets (2023). <https://github.com/alexlitel/congresstweets>.
- Converse, P. E. (1964). The nature of belief systems in mass publics (1964). *Critical review* 18(1-3), 1–74.
- Cormack, L. (2025). DCinbox: Official e-newsletters from every member of congress. <https://www.dcinbox.com/about/>.
- Costa, M. (2021). Ideology, not affect: What americans want from political representation. *American Journal of Political Science* 65(2), 342–358.

- Demszky, D., N. Garg, R. Voigt, J. Zou, M. Gentzkow, J. Shapiro, and D. Jurafsky (2019). Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. *arXiv preprint arXiv:1904.01596*.
- Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Di Tella, R., R. Kotti, C. Le Pennec, and V. Pons (2023). Keep your enemies closer: strategic platform adjustments during us and french elections. Technical report, National Bureau of Economic Research Cambridge, MA.
- Diermeier, D. and C. Li (2019). Partisan affect and elite polarization. *American Political Science Review* 113(1), 277–281.
- Druckman, J. N. and J. Levy (2022). Affective polarization in the american public. In *Handbook on politics and public opinion*, pp. 257–270. Edward Elgar Publishing.
- Eliasz, K. and R. Spiegler (2020). A model of competing narratives. *American Economic Review* 110(12), 3786–3816.
- Eliasz, K., R. Spiegler, and S. Galperti (2023). False Narratives and Political Mobilization. CEPR Discussion Papers 17832, C.E.P.R. Discussion Papers.
- Enke, B. (2020). Moral values and voting. *Journal of Political Economy* 128(10), 3679–3729.
- Fujiwara, T., K. Müller, and C. Schwarz (2024). The effect of social media on elections: Evidence from the united states. *Journal of the European Economic Association* 22(3), 1495–1539.
- Gehring, K. and M. Grigoletto (2023). Analyzing climate change policy narratives with the character-role narrative framework. Technical report, CESifo Working Paper.
- Gennaro, G. and E. Ash (2022). Emotion and reason in political language. *The Economic Journal* 132(643), 1037–1059.
- Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature* 57(3), 535–574.
- Gentzkow, M. and J. M. Shapiro (2010). What drives media slant? evidence from us daily newspapers. *Econometrica* 78(1), 35–71.
- Goetzmann, W. N., D. Kim, and R. J. Shiller (2022). Crash narratives. Technical report, National Bureau of Economic Research.

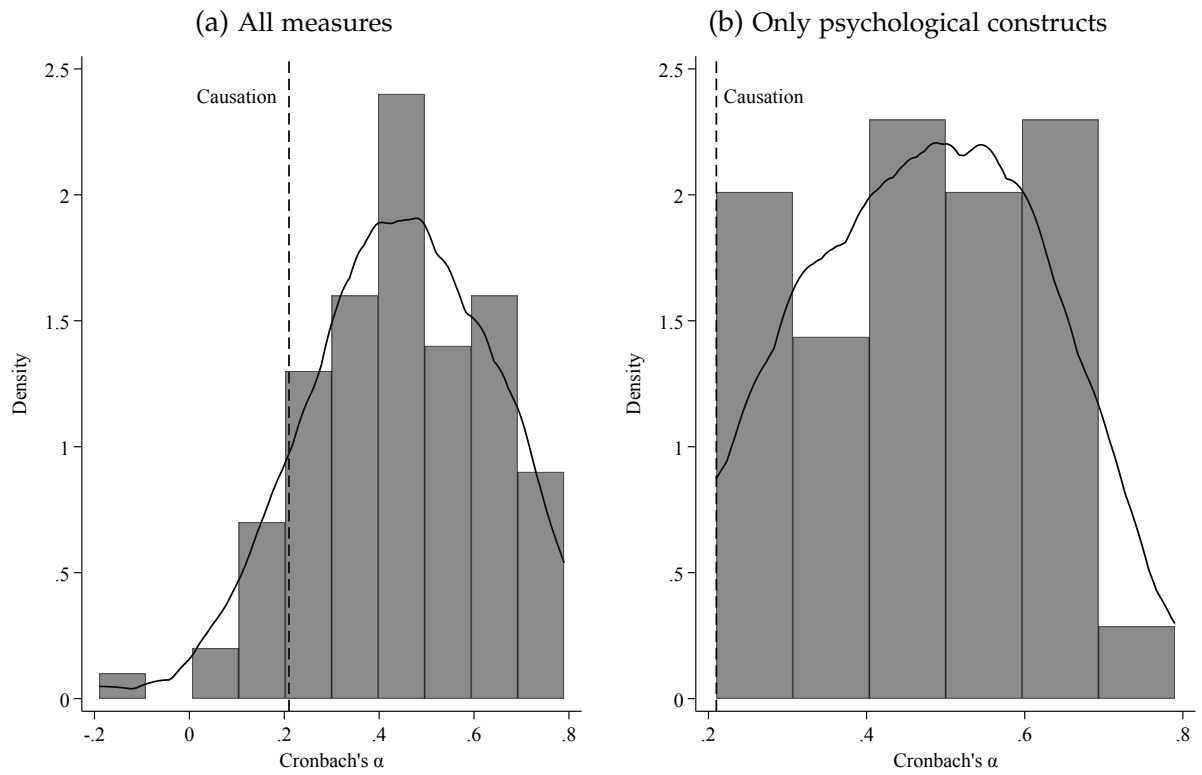
- Graeber, T., C. Roth, and C. Schesch (2024). Explanations. Technical report, CESifo Working Paper.
- Graeber, T., C. Roth, and F. Zimmermann (2024). Stories, statistics, and memory. *The Quarterly Journal of Economics* 139(4), 2181–2225.
- Graham, J., J. Haidt, and B. A. Nosek (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology* 96(5), 1029.
- Halberstam, Y. and B. Knight (2016). Homophily, group size, and the diffusion of political information in social networks: Evidence from twitter. *Journal of public economics* 143, 73–88.
- Healy, A. and N. Malhotra (2013). Retrospective voting reconsidered. *Annual review of political science* 16, 285–306.
- Hibbing, J. R., K. B. Smith, and J. R. Alford (2014). Differences in negativity bias underlie variations in political ideology. *Behavioral and brain sciences* 37(3), 297–307.
- Hirshleifer, D. (2020). Presidential address: Social transmission bias in economics and finance. *The Journal of Finance* 75(4), 1779–1831.
- Hood, C. (2010). The blame game. In *The Blame Game*. Princeton University Press.
- Hüning, H., L. Mechtenberg, and S. Wang (2022). Using arguments to persuade: Experimental evidence. Available at SSRN 4244989.
- Hutto, C. and E. Gilbert (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Volume 8, pp. 216–225.
- Ispano, A. (2025). The perils of a coherent narrative. *Economic Theory*, 1–15.
- Iyengar, S., Y. Lelkes, M. Levendusky, N. Malhotra, and S. J. Westwood (2019). The origins and consequences of affective polarization in the united states. *Annual review of political science* 22(1), 129–146.
- Iyengar, S., G. Sood, and Y. Lelkes (2012). Affect, not ideology: A social identity perspective on polarization. *Public opinion quarterly* 76(3), 405–431.
- Jain, A. (2023). Informing agents amidst biased narratives. Technical report, Mimeo.
- Kendall, C. W. and C. Charles (2023). Causal narratives. Technical report, National Bureau of Economic Research.

- Lau, R. R. (1985). Two explanations for negativity effects in political behavior. *American journal of political science*, 119–138.
- Lau, R. R. and I. B. Rovner (2009). Negative campaigning. *Annual review of political science* 12(1), 285–306.
- Lazarus, R. S. (1991). Cognition and motivation in emotion. *American psychologist* 46(4), 352.
- Liu, Y. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* 364.
- Lombrozo, T. (2012). Explanation and abductive inference. *The Oxford handbook of thinking and reasoning*, 260.
- Loureiro, D., F. Barbieri, L. Neves, L. E. Anke, and J. Camacho-Collados (2022). Timelms: Diachronic language models from twitter. *arXiv preprint arXiv:2202.03829*.
- Macaulay, A. and W. Song (2023). Narrative-driven fluctuations in sentiment: Evidence linking traditional and social media. Technical report, Bank of Canada Ottawa.
- Malhotra, N. and A. G. Kuo (2008). Attributing blame: The public’s response to hurricane katrina. *The Journal of Politics* 70(1), 120–135.
- Malle, B. F., S. Guglielmo, and A. E. Monroe (2014). A theory of blame. *Psychological Inquiry* 25(2), 147–186.
- Morag, D. and G. Loewenstein (2024). Narratives and valuations. *Management Science*.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. essay on principles. *Ann. Agricultural Sciences*, 1–51.
- Nguyen, D. Q., T. Vu, and A. T. Nguyen (2020). Bertweet: A pre-trained language model for english tweets. *arXiv preprint arXiv:2005.10200*.
- Pennebaker, J. W., R. L. Boyd, K. Jordan, and M. Blackburn (2022). Liwc 22. <https://www.liwc.app/>. Linguistic Inquiry and Word Count 22. Accessed: 2025-03-22.
- Piazza, J. and P. Sousa (2014). Religiosity, political orientation, and consequentialist moral thinking. *Social Psychological and Personality Science* 5(3), 334–342.
- Reimers, N. and I. Gurevych (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

- Rozin, P. and E. B. Royzman (2001). Negativity bias, negativity dominance, and contagion. *Personality and social psychology review* 5(4), 296–320.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and non-randomized studies. *Journal of educational Psychology* 66(5), 688.
- Schwartzstein, J. and A. Sunderam (2021). Using models to persuade. *American Economic Review* 111(1), 276–323.
- Shapiro, A. H., M. Sudhof, and D. J. Wilson (2022). Measuring news sentiment. *Journal of econometrics* 228(2), 221–243.
- Shiller, R. J. (2017). Narrative economics. *American economic review* 107(4), 967–1004.
- Sloman, S. A. and D. Lagnado (2015). Causality in thought. *Annual review of psychology* 66, 223–247.
- Sood, G. and S. Iyengar (2016). Coming to dislike your opponents: The polarizing impact of political campaigns. *Available at SSRN 2840225*.
- Stone, D. A. (1989). Causal stories and the formation of policy agendas. *Political science quarterly* 104(2), 281–300.
- Taylor, D. M. and J. R. Doria (1981). Self-serving and group-serving bias in attribution. *The Journal of Social Psychology* 113(2), 201–211.
- Weaver, R. K. (1986). The politics of blame avoidance. *Journal of public policy* 6(4), 371–398.
- West, D. M. (2018). *Air Wars: Television Advertising and Social Media in Election Campaigns, 1952-2016* (7th ed.). CQ Press.
- Westwood, S., Y. Lelkes, and M. Wetzel (2024). America’s political pulse: Elected officials.

A Additional Figures and Tables

Figure A1: LIWC Cronbach's α



Notes: Panel (a) plots the density of Cronbach's α for each 2022 LIWC subcategory, while Panel (b) restricts only to psychological constructs.

Table A1: Descriptive Statistics

	<i>Panel A) Tweet Level</i>			<i>Panel B) Politician Level</i>		
	All	Democrats	Republicans	All	Democrats	Republicans
Female	0.300 (0.458)	0.401 (0.490)	0.165 (0.371)	0.243 (0.429)	0.378 (0.486)	0.132 (0.339)
Age	57.973 (11.572)	59.058 (12.075)	56.509 (10.682)	57.468 (11.767)	58.499 (12.649)	56.617 (10.926)
Black	0.080 (0.272)	0.124 (0.330)	0.021 (0.142)	0.079 (0.270)	0.147 (0.355)	0.022 (0.148)
Bachelor	0.329 (0.470)	0.290 (0.454)	0.381 (0.486)	0.330 (0.470)	0.290 (0.454)	0.363 (0.481)
Master or Higher	0.623 (0.485)	0.676 (0.468)	0.551 (0.497)	0.586 (0.493)	0.654 (0.476)	0.529 (0.500)
Republican	0.426 (0.494)	0 (0)	1 (0)	0.548 (0.498)	0 (0)	1 (0)
Nominate Score	-0.005 (0.467)	-0.387 (0.126)	0.511 (0.164)	0.109 (0.455)	-0.366 (0.122)	0.500 (0.162)
Representative	0.804 (0.397)	0.814 (0.389)	0.790 (0.407)	0.857 (0.351)	0.857 (0.350)	0.856 (0.351)
Senator				0.118 (0.323)	0.123 (0.329)	0.114 (0.318)
Both Representative and Senator				0.026 (0.158)	0.020 (0.139)	0.030 (0.172)
Number of Tweets Posted				4664.950 (4848.342)	5924.391 (5151.747)	3625.209 (4319.150)
Number of Accounts				1.988 (0.689)	2.042 (0.688)	1.943 (0.687)
Observations	4198455	2411227	1787228	900	407	493

Notes: Panel A presents statistics at the tweet level. Panel B presents statistics at the politician level. Averages and standard deviations in parentheses.

Figure A2: Comparison of our Measures with Westwood et al. (2024)



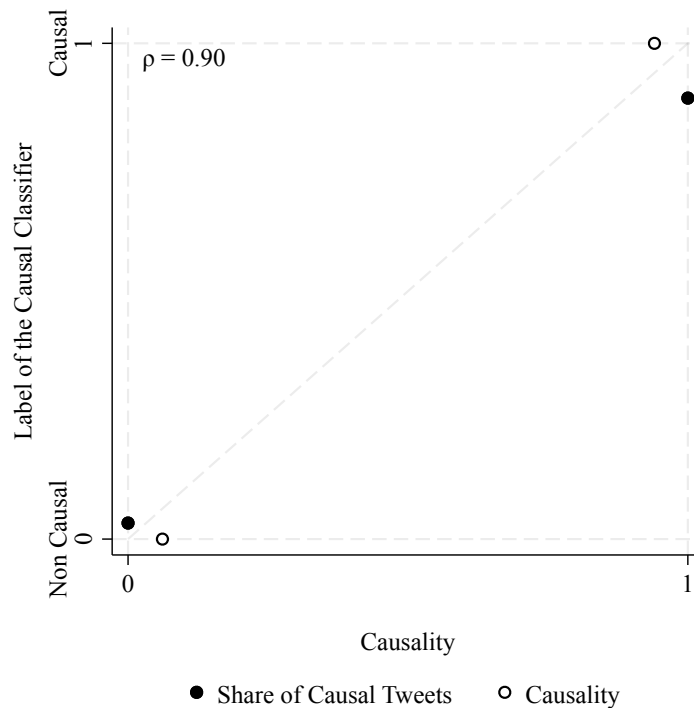


Figure A3: Classifier distributions




Notes: Panel (a) presents the distribution of probability that our classifier assigns to merit. Panel (b) does the same thing for blame, while Panel (c) for none.

Figure A4: Comparison of Causality with the Separate Causality Classifier



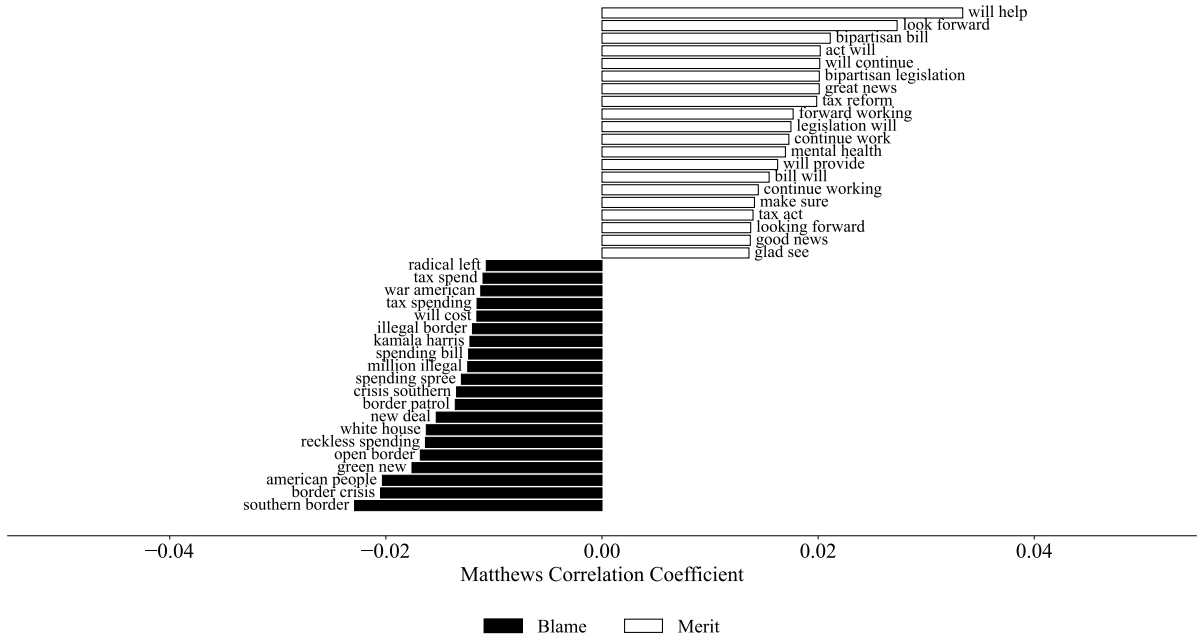
Notes: The figure presents that share of tweets classified as causal among the tweets with a Causality of 0 and 1, as well the share of tweets with a Causality of 1 among the tweets classified as non causal and causal. The correlation coefficient between the two binary vairable is displayed in the top left corner.

Figure A5: Merit and Blame Bigrams: Democrats



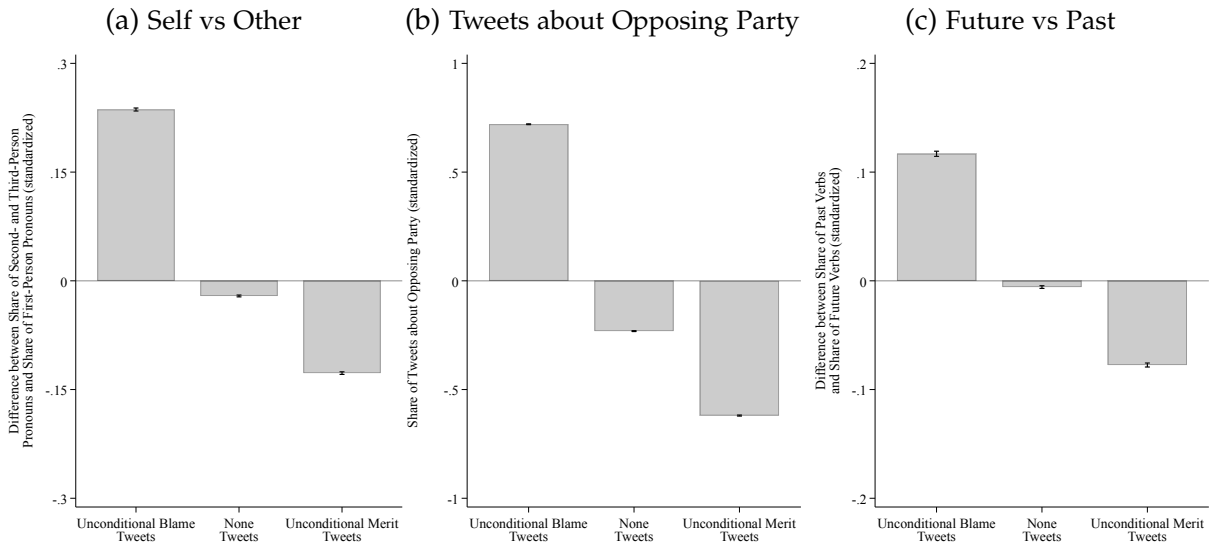
Notes: The figure presents 20 most distinctive bigrams of merit and blame tweets among Democrats, according to their Matthews Correlation Coefficient.

Figure A6: Merit and Blame Bigrams: Republicans



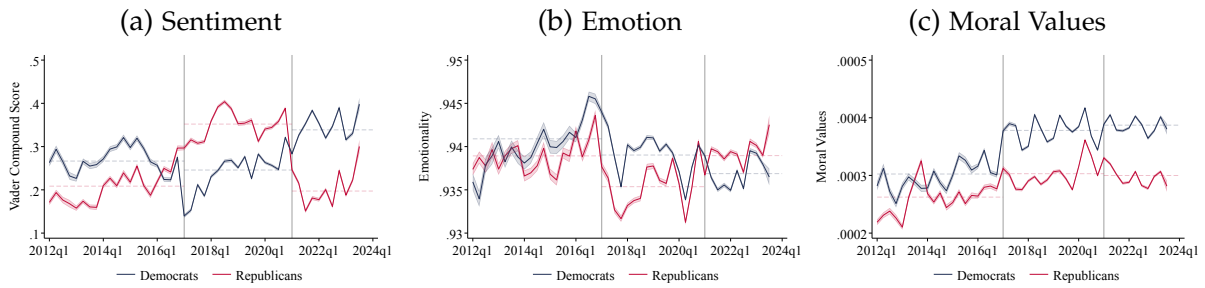
Notes: The figure presents 20 most distinctive bigrams of merit and blame tweets among Republicans, according to their Matthews Correlation Coefficient.

Figure A7: The ABC of Merit and Blame



Notes: Panel (a) presents the average difference between the share of words which are second- and third-person pronouns and the share of words which are first-person pronouns within blame, none, and merit tweets. Panel (b) presents the share of tweets the are about the opposing party within blame, none, and merit tweets. Panel (c) presents the average difference between the share of words which are past verbs and the share of words which are future verbs within blame, none, and merit tweets. In all panels the variable of interest is standardized.

Figure A8: Time trends in other text measures



Notes: Each panel reports, for the measure in its title, its quarterly evolution separately for Democrats and Republicans. Shaded areas denote 95 percent pointwise confidence intervals.

Table A2: Break Statistics from Sup-F Tests

	Quarter Break	Supremum F-Stat	Mean Before Break	Mean after Break	% Increase
Causality					
<i>Republicans</i>	2017q4	135.10	0.20	0.37	85.84
<i>Democrats</i>	2017q1	558.37	0.18	0.42	129.89
Sentiment					
<i>Republicans</i>	2016q2	22.75	0.20	0.29	47.19
<i>Democrats</i>	2020q4	47.04	0.25	0.34	34.99
Emotion					
<i>Republicans</i>	2020q4	6.47	0.94	0.94	0.23
<i>Democrats</i>	2020q1	17.26	0.94	0.94	-0.33
Moral Values × 1000					
<i>Republicans</i>	2016q2	36.11	0.26	0.30	16.56
<i>Democrats</i>	2017q1	189.02	0.30	0.38	28.15

Notes: Each panel reports results based on the `sbsingle` command in Stata, which estimates an unknown-date structural break for each time series separately. Averages before and after the break are computed on the entire sample.

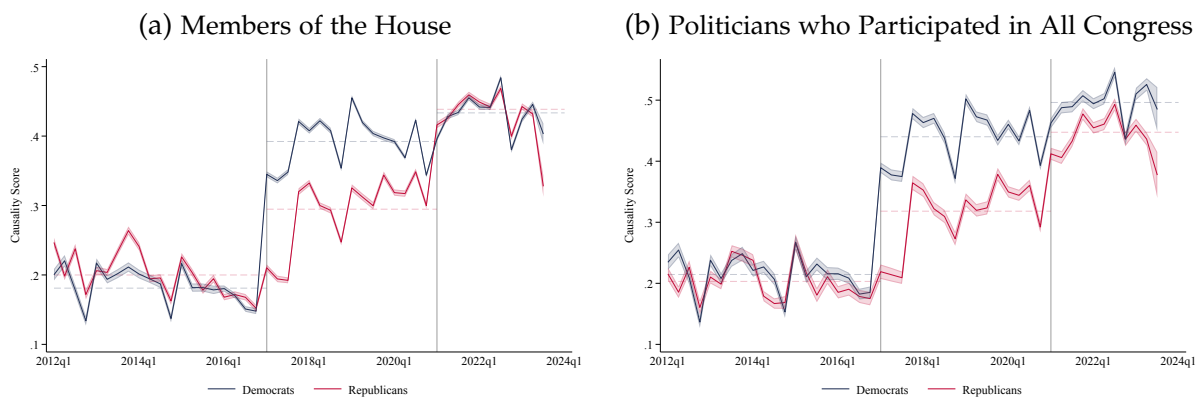
Table A3: Seemingly Unrelated Regressions

	Point Estimate	Lower Bound	Upper Bound
$\beta_{\text{Causality}} - \beta_{\text{Sentiment}}$			
<i>Republicans</i>	0.08	0.01	0.14
<i>Democrats</i>	0.21	0.17	0.25
$\beta_{\text{Causality}} - \beta_{\text{Emotion}}$			
<i>Republicans</i>	0.16	0.13	0.19
<i>Democrats</i>	0.24	0.21	0.26
$\beta_{\text{Causality}} - \beta_{\text{Moral values}}$			
<i>Republicans</i>	0.16	0.13	0.19
<i>Democrats</i>	0.23	0.21	0.26

Notes: Each panel reports results based on the `sureg` command in Stata, which estimates a SUR model with robust standard


errors. Linear combination tests are taken with the `lincom` command. Upper and lower bound refer to 95 percent confidence intervals.

Figure A9: Causality Evolution in Subsamples




Notes: The figure presents the average Causality at the quarter level separately for Democrats and Republicans, in Panel (a) restricting only to Members of the House and in Panel (b) to politicians who have been elected in all Congresses from 112 to 118. Shaded areas denote 95 percent pointwise confidence intervals.

Figure A10: Event Studies: Causality



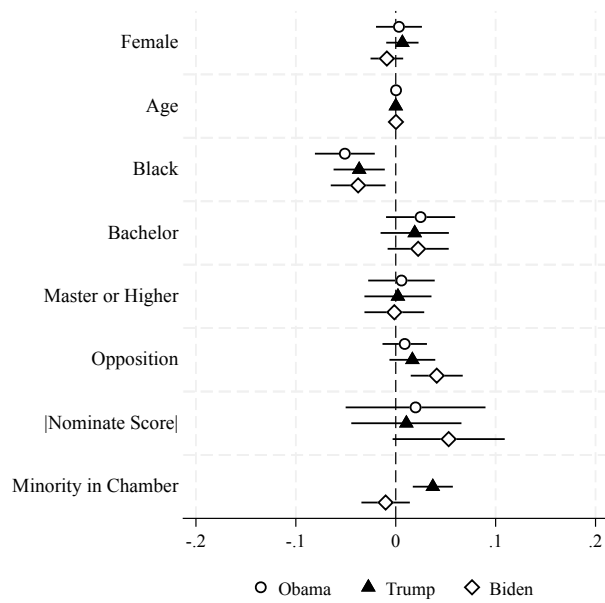
Notes: The figure presents the estimates of two event studies at the politician level for Causality, with base level at 2012q1, computed separately for Democrats and Republicans. Shaded areas denote 95 confidence intervals with standard errors clustered at the politician level.

Figure A11: Individual Shifts



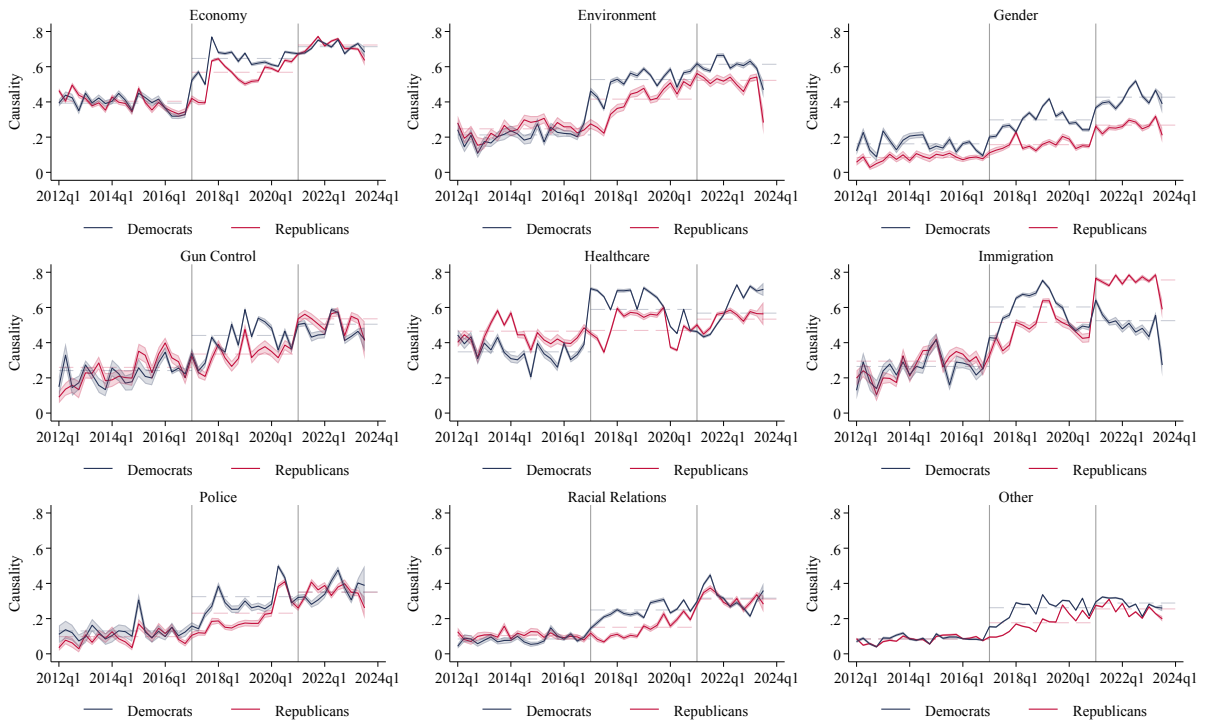
Notes: The coefficient reports the distribution of β s from the regression $\text{Causality}_{i,t} = \alpha + \beta \mathbb{1}\{t \geq \tau\} + \varepsilon_{i,t}$ where $t \geq \tau$ whenever tweet i was posted after 2016, estimated separately for each politician with robust standard errors.

Figure A12: Politician Correlates of Causality




Notes: The coefficients presents the estimates from three regressions at the politician level of Causality over the listed variables and the text-level measures of Figure 1, computed separately for each presidency. Bars denote 95 percent confidence interval with robust standard errors.

Figure A13: Time Trend: Causality by Topics



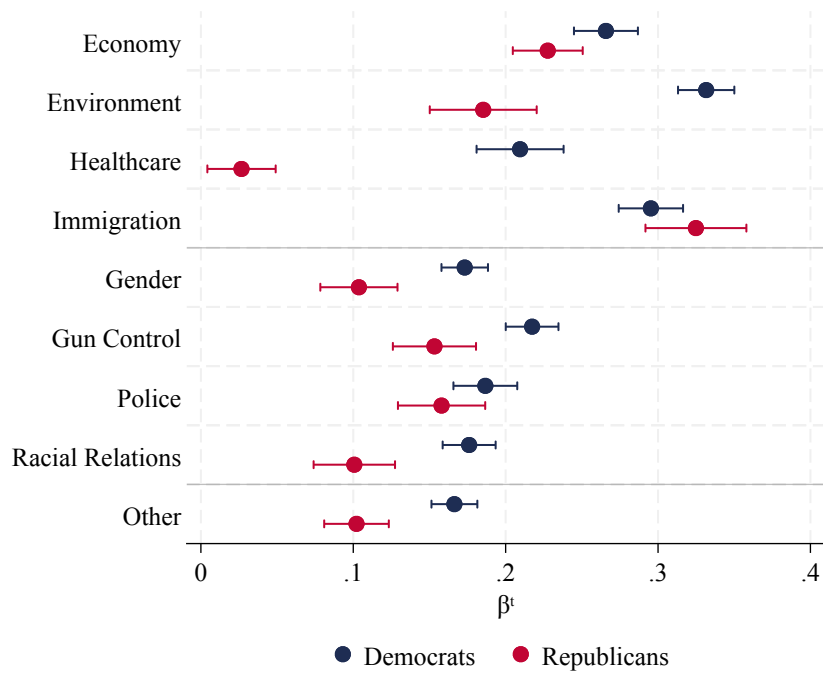
Notes: The figure presents the average Causality at the year level separately for Democrats and Republicans separately for each topics. Shaded areas denote 95 percent pointwise confidence intervals.

Figure A14: Topic Distribution by Party




Notes: The figure presents the share of tweets within each topic in each of the three presidency, separately for Democrats and Republicans.

Figure A15: Within Topic Shift in Causality




Notes: The figure presents the estimates of Equation 3. Bars denote 95 percent confidence intervals with standard errors clustered at the politician level.

Figure A16: Causality over Topics, by Party



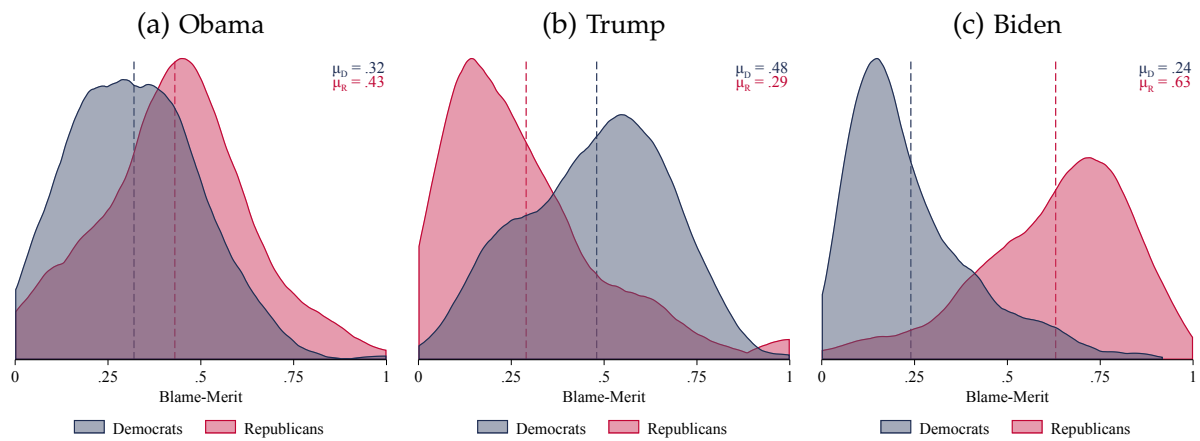
Notes: The figure presents the average Causality for each topic, separately for Democrats and Republicans. Bars denote 95 percent confidence intervals.

Figure A17: Politician Correlates of Blame



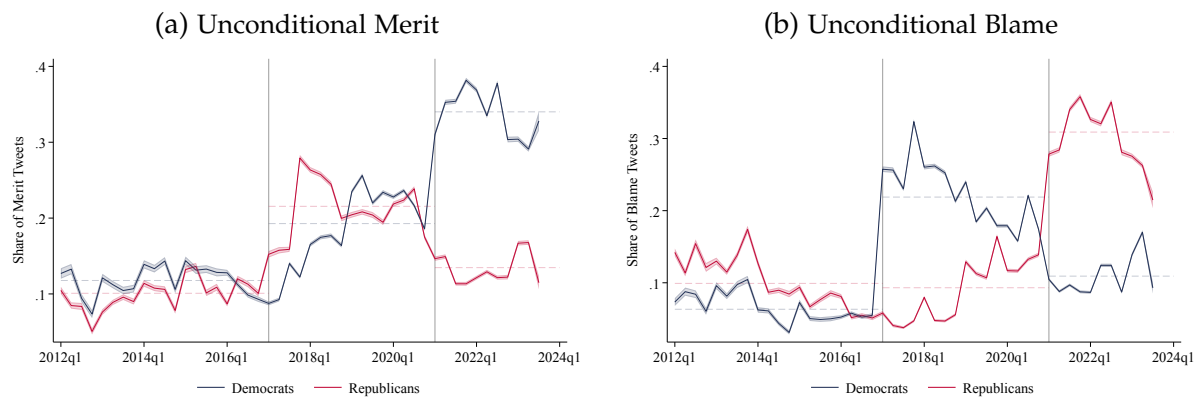
Notes: The coefficients presents the estimates from three regressions at the politician level of Causality over the listed variables and the text-level measures of Figure 1, computed separately for each presidency. Bars denote 95 percent confidence interval with robust standard errors.

Figure A18: Blame Densities



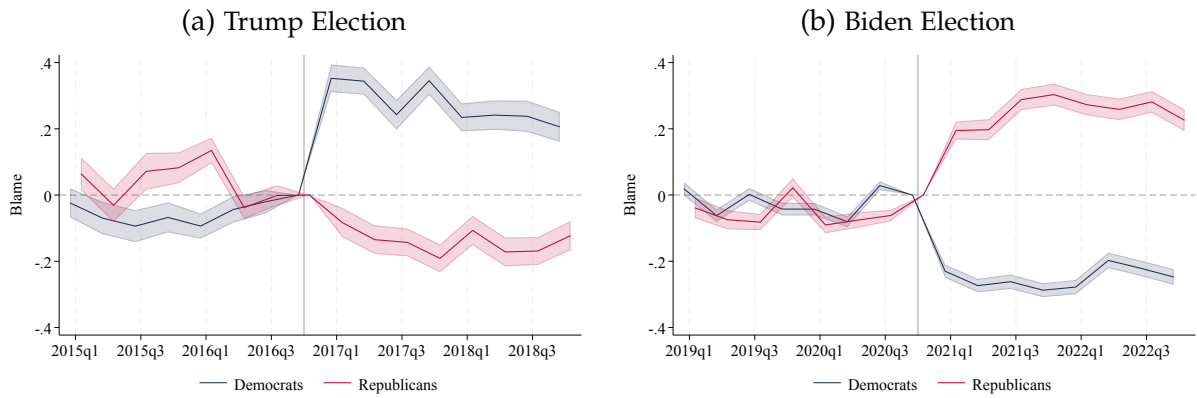
Notes: Panel (a) presents the density of Blame during the Obama presidency. Panel (b) presents the density of Blame during the Trump presidency. Panel (c) presents the density of Blame during the Biden presidency. In all panels the blue (red) dashed line is the average level of Blame among Democrats (Republicans), with the corresponding value reported in the top right corner in blue (red).

Figure A19: The Evolution of Unconditional Merit and Blame



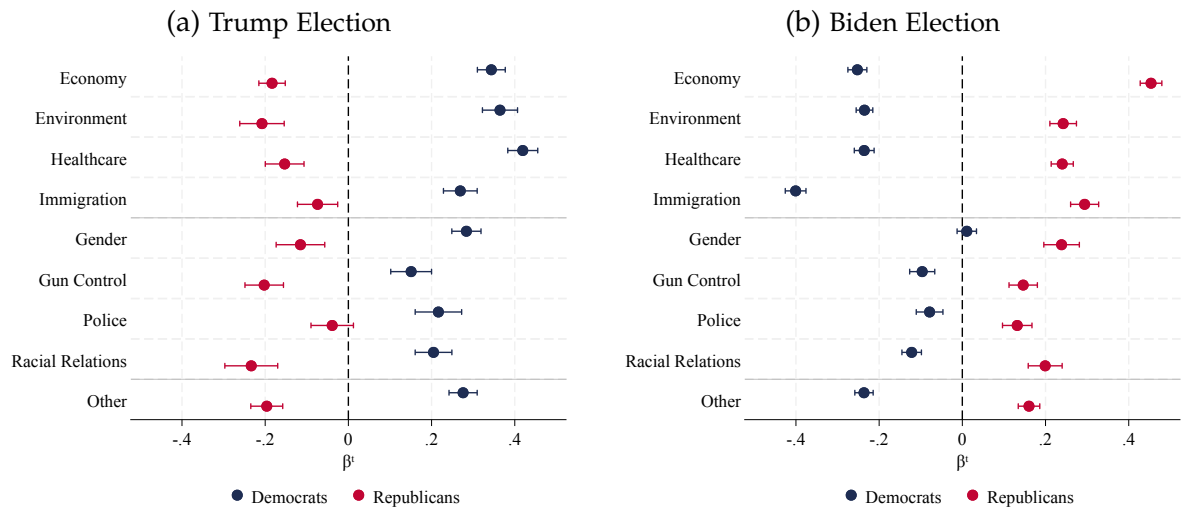
Notes: Each panel reports the share of merit (blame) tweets over all tweets averaged at the quarterly level, separately for Democrats and Republicans. Shaded areas denote 95 percent pointwise confidence intervals.

Figure A20: Blame Shifts around Elections




Notes: Both panels report coefficients from Equation 2 where, instead of causality, the outcome is blame. For Panel (a), the sample is restricted between 2015 and 2018 (both included), while for Panel (b) the restriction is between 2019 and 2022. Bars denote 95 percent confidence intervals with standard errors clustered at the politician level.

Figure A21: Blame Shifts around Elections, across Topics



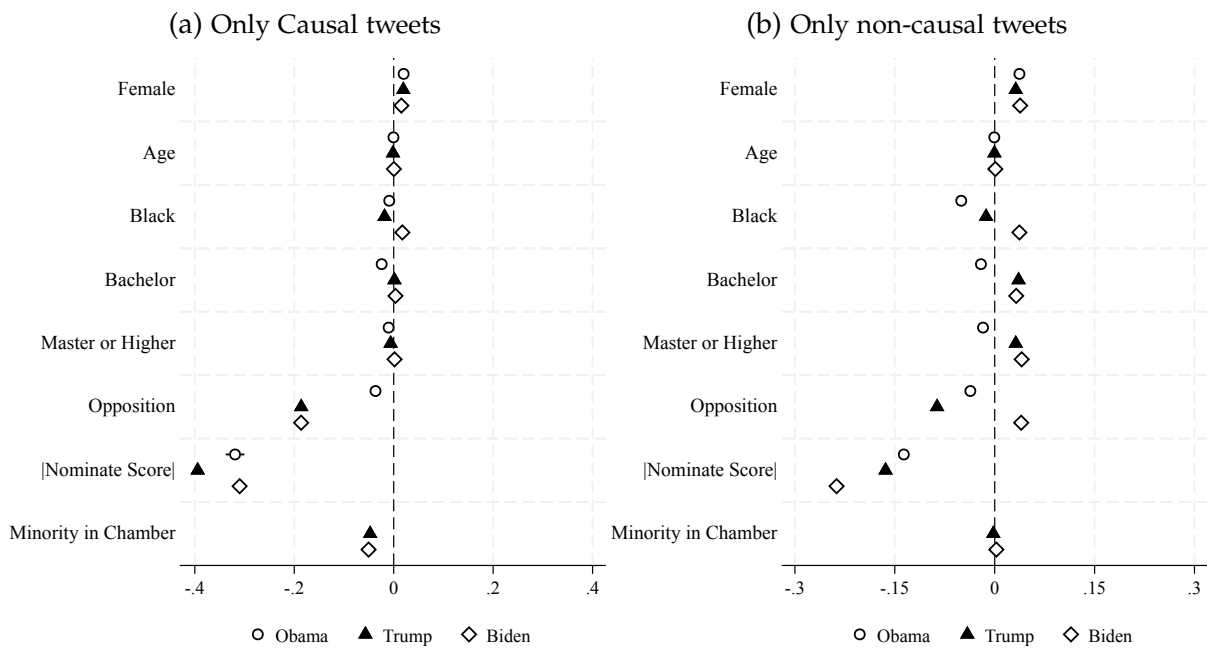
Notes: Both panels report coefficients from Equation 3 where, instead of causality, the outcome is blame. For Panel (a), the sample is restricted between 2015 and 2018 (both included), while for Panel (b) the restriction is between 2019 and 2022. Bars denote 95 percent confidence intervals with standard errors clustered at the politician level.

Figure A22: Causality and Sentiment




Notes: The figure presents the results from regression 4. Bars denote 95 percent confidence intervals.

Figure A23: Causality and Sentiment Separately



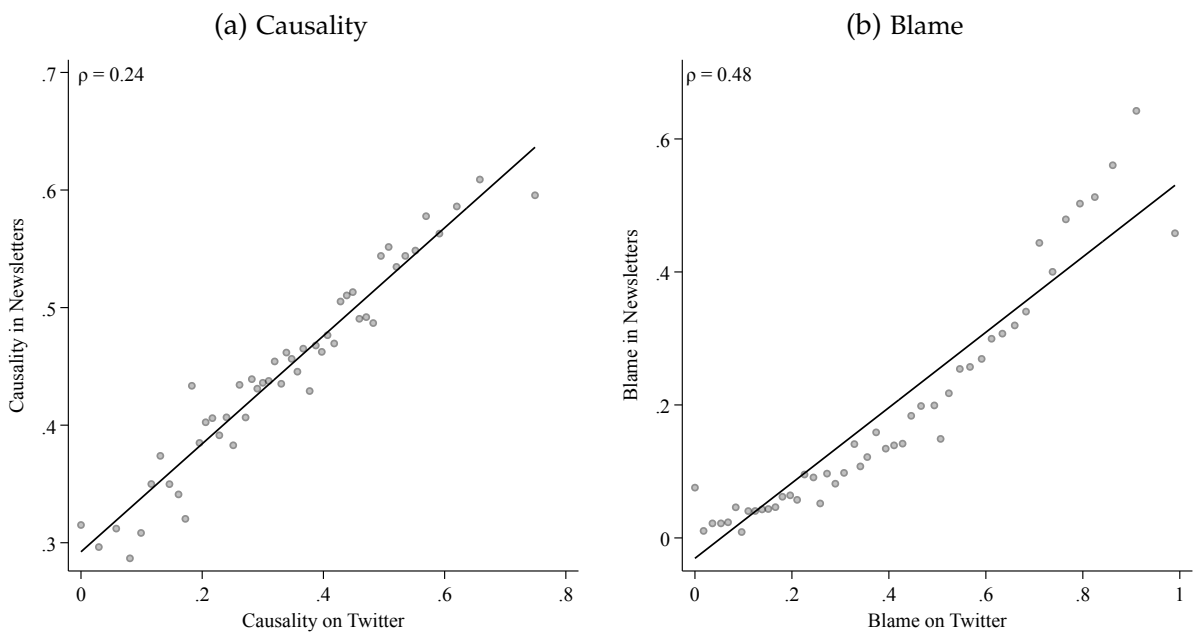
Notes: Bars denote 95 percent confidence intervals with standard errors clustered at the politician level.

Figure A24: Causality in Newsletters



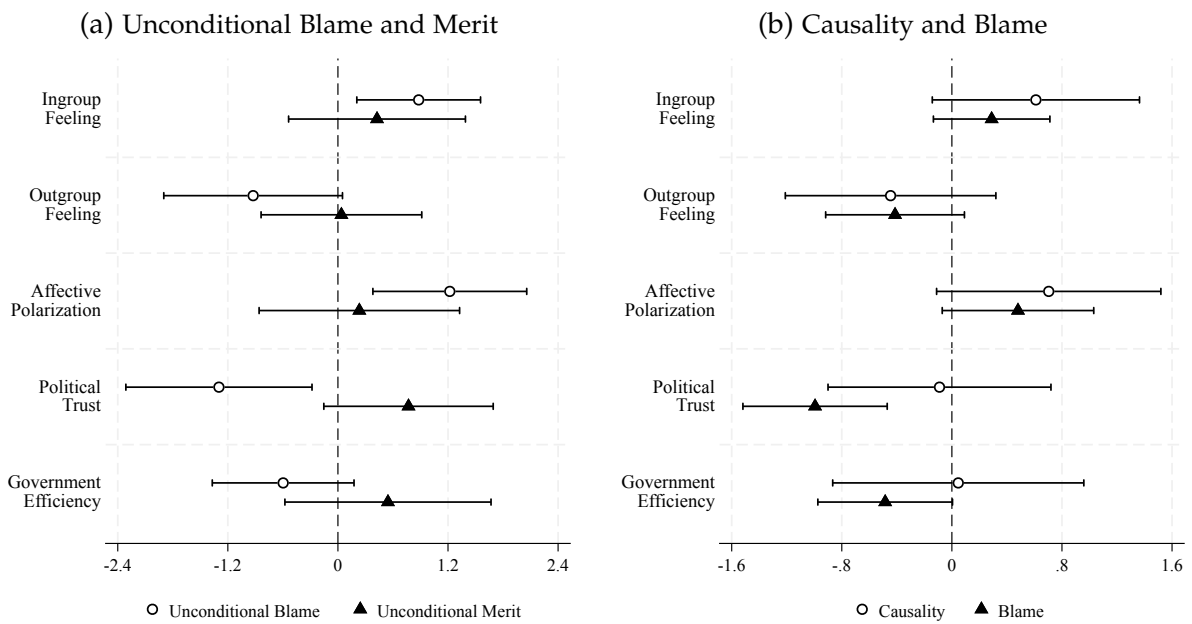
Notes: the y-axis reports at quarterly level the share of causal sentences for each party. Shaded areas denote 95 percent pointwise confidence intervals.

Figure A25: Causality and Blame: Twitter vs. Newsletters



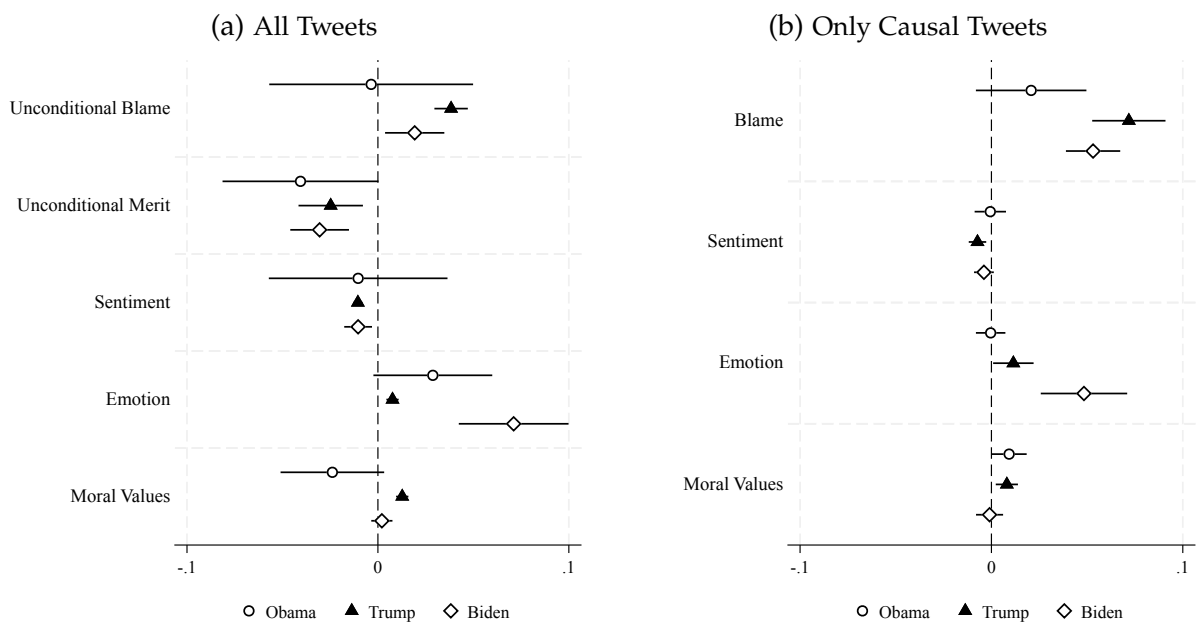
Notes: In Panel (a) we plot on the y-axis the share of causal sentences for each politician in each quarter, while on the x-axis the share of causal tweets for the same politician in the same quarter. Panel (b) does the same for the share of blame sentences among causal ones.

Figure A26: Correlation with People’s Beliefs and Attitudes from Westwood et al. (2024)



Notes: Panel (a) presents the estimates of five separate regression of the five listed outcome variables over the share of merit and blame tweets. Panel (b) presents the estimates of five separate regression of the five listed outcome variables over our indicated measures. In both panels, all the estimates are at the state-week levels and include week fixed effects. In both panels, bars denote 95 percent confidence intervals with standard errors clustered at the state level.

Figure A27: Presidents’ Retweets



Notes: All panels present regressions where the outcome variable is retweets, the regressors are those listed on the y-axis, and include week and topic FE. Errors are clustered at the week level, and regressions are estimated separately for each president. Both the outcome and the regressors are standardized. In both panels, bars denote 95 percent confidence intervals with standard errors clustered at the state level.

Table A4: Labeled Tweets

Tweet	Causal	Tone	Rhetorical Style
Biden has lost all credibility. No one believes your lies, Joe! URL	0	-1	None
Juan Williams just compared Obamacare to a sweater. Good analogy. It's something you don't want to get, but have to accept when given.	0	-1	None
Joe Biden and the Democrats terrible policies have wreaked havoc on this country.	1	-1	Blame
#Trumpcare is fundamentally flawed. Higher costs, less coverage, fewer protectionsthats GOPs gift to American people. #ProtectOurCare URL	1	-1	Blame
I couldnt support final passage of todays approps package but Im pleased about the inclusion of my HBCU amendment URL	0	1	None
Universal congrats to the scientists at @OregonState for their work helping Insight make a successful #MarsLanding: URL	0	1	None
The Protect Medical Innovation Act will boost American innovation and manufacturing, and it will encourage medical research and development that make a real difference in peoples lives. URL	1	1	Merit
Now that the Inflation Reduction Act is law it will not only lower prescription drug prices but save lives. Thank you @HenryFordHealth for your support. URL	1	1	Merit

Table A5: Democrats Tweets Classified as Merit and Blame

Panel (a) Merit Tweets

Proud to cosponsor @NydiaVelazquez's Public Housing Emergency Response Act which would invest \$70B in public housing including \$32B for @NYCHA. Our public housing crisis must be addressed & this bill is a bold approach to doing that. Residents deserve to live in safe conditions!

I applaud @POTUS for setting our next offshore wind target. With a new infusion of investments from my offshore wind manufacturing tax credit in the Inflation Reduction Act, the U.S. can and will deploy 15 GW of floating offshore wind by 2035 all while creating good union jobs. URL QT @ginamccarthy46 Today we're announcing actions to advance *floating* offshore wind platformskey to harnessing the potential of deep waters along the West Coast, Gulf of Maine, & more. Part of @POTUS' plan for American jobs and leadership on new clean energy technologies! URL

The Bipartisan Infrastructure Law is putting our economy on track to thrive & investing in communities that have too often been left behind. With over \$2 million recently announced, we take a major step toward redeveloping Baltimores Bridge to Nowhere. URL

Happy to be joining @HouseDemocrats to help Americas workers access better paying jobs. The Workforce Innovation and Opportunity Act connects employers with qualified candidates, lowers costs for families and increases supplies. Democrats are #InvestingInWorkers. URL

The #BuildBackBetterAct provides much needed funds to @TheJusticeDept to help reduce community violence & fund proven intervention programs. Im proud to advocate for legislation to break cycles of violence in communities, saving American lives & taxpayer dollars.

Panel (b) Blame Tweets

We warned when the GOP passed tax cuts for the rich that it would explode deficits. It did. We warned that the GOP would use those deficits to come after Social Security and Medicare. They are. URL

The party of NO, #ILGOP in particular, plays political games in ignoring the implications on our economy, on jobs, Social Security checks. Republicans raised the debt ceiling 3 times under Trump's thumb. They are playing politics with people's lives. URL

GOPs reckless health care strategy is already destabilizing #healthcare markets and forcing premiums to rise. and kids health care at risk, pensions at risk, and the fight against opioids at risk

Real wages today are lower than they were in 1973. Thats not a sign of a healthy economy, its a sign that working people today are worse off than they were 45 years ago, and the GOP tax cuts have done nothing to address that issue.

Table A6: Republicans Tweets Classified as Merit and Blame

Panel (a) Merit Tweets

I'm an original sponsor of the Nuclear Energy Leadership Act w/ @RepElaineLuria to encourage further dev of advanced nuclear energy programs. Such programs will create high-quality jobs, strengthen natl security, reduce foreign energy dependence and promote emissions-free energy.

Today @realDonaldTrump showed his commitment to supporting American energy dominance. The @EPA's rule will bolster our nations energy independence by lowering energy costs, spurring job growth and promoting economic development in our communities. URL

Glad to hear @realDonaldTrump has signed legislation adding \$320 billion to the #PaycheckProtectionProgram, ramping up testing capability and providing more funding for health care providers. AR will benefit from this measure to protect public health and save businesses & jobs. URL URL

Earlier this week, I introduced #CARA2 to increase funding levels for programs we know work and implements additional policy reforms that will make a real difference in combatting the #opioidcrisis. URL

@TransportGOP are delivering on our promise of fixing supply chain holes and building a stronger economy. Currently we are marking up a package of bills that will remove barriers, increase efficiency, and target infrastructure investment. #SupplyChain

Panel (b) Blame Tweets

The supply chain and inflation crises are not a high class problem like @WHCOS claims. As Dems look to pour trillions into the economy and spike inflation further, they must understand that actions have consequences that will be felt by every American URL

Our country is facing soaring inflation thanks to Democrats spending spree, and whats @POTUS solution? Spend MORE money. @JoeBiden and @SenateDems are TOTALLY out of touch with reality. Inflation is still wiping out wage growth, all while Democrats' reckless spending spree makes matters worse. URL

Top border officials told Biden that if he unraveled Trump's policies and pushed for open borders that a major crisis would occur. He didnt listen. Now everyone is suffering - Americans and migrants alike. URL

April saw the highest number of migrants ever recorded. Next week, @JoeBiden will reverse another commonsense border policy that will only make this crisis worse. Biden needs to wake up and face reality. URL
