

4-11-2024

Framing contestation and public influence on policymakers: evidence from US artificial intelligence policy discourse

Daniel S. Schiff
Purdue University, dschiff@purdue.edu

Follow this and additional works at: <https://docs.lib.purdue.edu/fund>



Part of the [Political Science Commons](#)


Recommended Citation

Schiff, Daniel S., "Framing contestation and public influence on policymakers: evidence from US artificial intelligence policy discourse" (2024). *Purdue University Libraries Open Access Publishing Fund*. Paper 199.

<http://dx.doi.org/10.1093/polsoc/puae007>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries.
Please contact epubs@purdue.edu for additional information.

Framing contestation and public influence on policymakers: evidence from US artificial intelligence policy discourse

Daniel S. Schiff 

Department of Political Science, Purdue University, 100 North University Street, West Lafayette, IN 47907, USA

Corresponding author: Daniel S. Schiff, Department of Political Science, Purdue University; 100 North University Street, West Lafayette, IN 47907, USA. Email: dschiff@purdue.edu

Abstract

As artificial intelligence (AI) policy has begun to take shape in recent years, policy actors have worked to influence policymakers by strategically promoting issue frames that define the problems and solutions policymakers should attend to. Three such issue frames are especially prominent, surrounding AI's economic, geopolitical, and ethical dimensions. Relatedly, while technology policy is traditionally expert-dominated, new governance paradigms are encouraging increased public participation along with heightened attention to social and ethical dimensions of technology. This study aims to provide insight into whether members of the public and the issue frames they employ shape—or fail to shape—policymaker agendas, particularly for highly contested and technical policy domains. To assess this question, the study draws on a dataset of approximately five million Twitter messages from members of the public related to AI, as well as corresponding AI messages from the 115th and 116th US Congresses. After using text analysis techniques to identify the prevalence of issue frames, the study applies autoregressive integrated moving average and vector autoregression modeling to determine whether issue frames used by the public appear to influence the subsequent messaging used by federal US policymakers. Results indicate that the public does lead policymaker attention to AI generally. However, the public does not have a special role in shaping attention to ethical implications of AI, as public influence occurs only when the public discusses AI's economic dimensions. Overall, the results suggest that calls for public engagement in AI policy may be underrealized and potentially circumscribed by strategic considerations.

Keywords: agenda-setting; artificial intelligence; issue framing; public participation; technology governance

Technology governance is at a crossroads. Given the role of advanced information technology in economic transformation, emerging technologies have become pillars of national innovation strategy (Soete, 2007). Simultaneously, recognition of the power accruing to industrial and national technology leaders has inspired arguably unprecedented calls for attention to technology's social and ethical risks. This is nowhere more true than with respect to artificial intelligence (AI), a strategically important general purpose technology often construed as the pillar of 21st century innovation but marked by substantial alarm about a wide array of potential social and ethical harms and situated in the middle of a global technology race (Ulnicane, 2023).

In the context of these tensions, a multisectoral consensus has emphasized the need for the public to exert agency in shaping the emerging AI policy agenda (Buhmann & Fieseler, 2023). This call for diverse

participation in what might otherwise be an expert-dominated policy domain aligns with newfound attention to notions of shared governance (Minkkinen et al., 2023; Pierre, 2000). Indeed, policy actors across the public, private, and non-governmental sectors agree that members of the public should play a critical role in weighing the benefits and risks of AI and informing ultimate decision-making (Crockett et al., 2021). In particular, the public is argued to have a special stake in informing government about their concerns with respect to social and ethical risks of AI. How AI governance unfolds, then, may also depend on how AI is understood—or framed—as a policy issue (Imbrie et al., 2021).

Yet, notwithstanding this apparent normative consensus, what is less clear is whether the public actually has meaningful opportunities to shape the AI policy agenda. In turn, this study seeks to answer the following in the context of the US, a global leader in AI research and development and key policy actor:

- Does the public shape policymaker attention to AI, despite the fact that technology policy is traditionally expert-dominated and not highly salient to the public?
- Given the public's special stake in social and ethical issues, is the public especially influential in shaping policymaker attention when AI is framed in terms of its social and ethical implications?

To assess these questions, datasets were created aimed at capturing attention to several issue frames related to AI policy as invoked by three groups of actors: the public, federal policymakers, and news media in the US, covering the period of 2017 through 2019.¹ Using quantitative text analysis techniques, this study extracted actor-specific time series datasets representing attention to three issue frames, reflecting potentially competing concerns about AI's prospects for innovation, social and ethical dimensions, and implications for geopolitical competition. Next, to analyze whether the public shapes policymaker attention to AI in general or only with respect to certain issues frames, this study applied autoregressive integrated moving average (ARIMA) and vector autoregression (VAR) modeling.

Results across a variety of specifications indicate that the public has a heightened but ultimately limited role in shaping policymaker attention with respect to the emerging AI policy agenda. Notably, public attention to AI *does* predict policymaker attention to AI, and no other relationship between members of the public, media, and policymakers exhibits this degree of influence. Yet this ostensible influence only occurs when the public discusses AI generally or with respect to its role in economic innovation. That is, *the public plays no special role in driving attention to AI's social and ethical implications*, potentially challenging assumptions about the public's priorities or the nature of policymakers' responsiveness to them.

This study contributes by advancing the use of text analysis and time series methods to assess questions in policy process theory and innovation policy, particularly related to the emergence of the US AI policy agenda. Through empirical study of *sub-issue* attention to issue frames and the role framing contestation plays in agenda-setting, this research advances the understanding of agenda-setting dynamics for strategic and emerging technology policy domains, and provides key insight into whether the public plays a meaningful role in technology governance. The results reveal that the public does matter, but perhaps not in the way expected or imagined.

Theoretical approach

New directions in technology governance

The governance of science and technology may be changing in the 21st century. Limitations of traditional science and technology policy, such as its failure to solve major societal challenges like inequality and climate change (Uunicane, 2016), have led scholars to propose new approaches to innovation policy altogether. For example, movements like transformative innovation policy, mission-oriented innovation policy, and Responsible Research and Innovation (Diercks et al., 2019; von Schomberg, 2013) deviate from a more economic, expert-driven, and firm-centered logic for innovation. While these approaches are evolving, remain contested, and skirt the line between normative and descriptive, they share certain key features and can be jointly understood to reflect a new aspirational paradigm (Hogan & Howlett, 2015) in technology policy.

¹ This analysis allows explicit attention to a single presidential administration, useful for analytical and methodological reasons. Yet, it leaves open important questions surrounding the evolution of AI policy, public participation, and issue framing in the 2020s which this paper cannot conclusively address.

In particular, such a paradigm emphasizes attention to social rather than purely economic goals for innovation (Kuhlmann & Rip, 2018). Along these lines, it acknowledges the possibility that innovation can lead to ethical and societal harms, such that a ubiquitous pro-innovation bias may be problematic. Further, it implies a more inclusive and participatory societal policy agenda, for which it is important to incorporate a broader array of actors beyond firms, academia, and government, such as civil society and especially the public (Warnke et al., 2016). An important question for the future of technology governance is therefore whether critical emerging policy domains like AI are in fact reflective of elements of this new paradigm, namely an increased reliance on public participation and a focus on social and ethical dimensions of technology.

Broadening participation in AI policy

The central question for this study is whether—and under what conditions—the public shapes the technology policy agenda, specifically in terms of public influence on federal policymaker discourse in US AI policy. Normatively speaking, there is a substantial multi-sector and international consensus on the value of public participation in AI governance (Ulnicane et al., 2021; Vesnic-Alujevic et al., 2020). Members of the public are argued to be deeply affected by the social transformations resulting from AI, implying that they have a special stake and that policymakers need to be responsive to public concerns (Stahl, 2021). For instance, controversies related to privacy violations or social media platforms and—more recently—the growth of generative AI may play a role in the public's increasing understanding of AI's implications. As such, members of the public may function as “citizen experts” (Fischer, 2000) given first-hand experience of AI-based harms, such as faulty automated decision systems used by government, labor displacement resulting from automation, and algorithmic surveillance.

Reflection on the importance of public participation in policy extends from a long history of political thought, echoed in broader 21st century approaches to governance rather than top-down government alone (Pierre, 2000; Rowe & Frewer, 2000). Yet, in the context of science and technology policy, the perceived heightened risks of technological impacts and the gap between subject expert and public understanding (Dryzek & Pickering, 2017) have led to urgent calls for public participation in technology governance generally (Macnaghten and Chilvers 2014; Stirling, 2008), and in AI policy specifically (König and Wenzelburger 2021; Stark et al., 2021). In particular, actors in the public, private, and non-governmental sectors have called for increased public scrutiny, public discourse, and participation of specialized mini-publics to inform public and private sector AI governance (Buhmann & Fieseler, 2023). For example, Ouchchy et al. (2020) find that “encouraging public involvement” is the single most common recommendation made in media coverage of AI.

Yet, evidence on the willingness, ability, and capacity of policymakers to engage the public and the ultimate influence of public participation on policy agendas is decidedly mixed (Abelson & Gauvin, 2006; Lodge & Wegrich, 2015). For example, the Multiple Streams Framework (MSF) considers public opinion (or national mood) to be important in shaping the policy agenda (Kingdon, 1984). Yet, Barberá et al. (2019) find in a recent study of 100 policy topics that the general public does not lead issue priorities for Congress. In the AI space in particular, some have argued that calls for public participation may represent “cheap talk” and emphasize superficial or self-serving aims, enabling what Sloane et al. (2020) term “participation-washing.” Policymakers or public and private interest groups might mischaracterize public attitudes, promote only public audiences that already share one's views, or express interest in public participation merely as a signaling strategy to avoid more extensive interventions. Finally, it remains unclear whether policymakers know how to effectively solicit public opinion even when the intention to do so is genuine.

This study contributes to scholarly efforts in policy process theory and innovation policy to understand the role of the public in agenda-setting by considering whether public attention to AI (i.e., within the public or systemic agenda) is echoed by policymakers (i.e., in the institutional or Congressional agenda) for emerging technology policy domains.² If the calls for public participation in AI are bearing out in practice, we should expect to see increased attention by policymakers to the concerns of the general public. Thus, a first hypothesis is that:

² A note of causal caution is prudent. The methods in this study may only be able to nod towards causation in a limited fashion; in that sense, public-policy issue attention correlation can be understood in terms of “issue congruence” (Jones & Baumgartner, 2004) rather than strict causation.

Public agenda-setting hypothesis

Issue attention by members of the public to AI policy generally will predict issue attention by policymakers.

However, it may be the case that public attention only predicts policymaker attention when AI is conceived of in certain ways, i.e., with respect to specific *issue frames*. For example, policymakers may feel more pressure to be responsive to public attention with respect to AI's ethical and social issues, but rely more on the expertise of military or economic specialists in closed-door settings when international competition and technological leadership are at stake.

Contested issue frames in AI agenda-setting discourse

Issue frames attempt to capture an ongoing discourse and interpret events and indicators in a way that packages multifaceted issues in terms of a more simplified essence (Chong & Druckman, 2007; Gamson & Modigliani, 1989). In this way, frames structure the underlying categories associated with a topic, strategically emphasizing or excluding certain elements, in order to constrain or expand reasoning in preferred directions (Baumgartner & Jones, 1993; Sharp, 1994). Powerful and effective frames consequently may direct both public and elite attention, legitimizing certain policy narratives and bringing particular policy solutions into the mainstream (Goddard & Krebs, 2015). At one level then, the impact of frames on agenda formation is straightforward: When specific issue frames become dominant, there is a greater chance that policymakers will respond to these signals by putting these issues and solutions to them on the institutional agenda (Peters & Hogwood, 1985).

In the context of AI policy, this study examines three issue frames prominent in AI discourse, exemplified in Figure 1. The first frame is the “innovation frame.” In line with traditional technology and innovation policy, this frame emphasizes AI's potential for economic transformation in terms of fostering industrial productivity, economic efficiency gains, high-tech growth acceleration, entrepreneurial activity, and national GDP expansion (Edler & Fagerberg, 2017; Soete, 2007). Advocates of this issue frame for AI typically emphasize the need for supply-side reforms, such as education and training efforts to increase the size of the AI workforce, increased research and development funding, and regional and national industry–government–university coordination (Fischer et al., 2021; Smits & Kuhlmann, 2004). For example, the United States Innovation and Competition Act, American AI Initiative, and National AI Research and Development Strategic Plan deeply embody this thinking.

A second “often invoked” frame emphasizes social and ethical implications of AI—an “ethics frame” (Ulnicane et al., 2020, 2). A sizable body of work by civil society, epistemic communities of academic experts and practitioners, and the private sector has reviewed numerous ethical problems associated with AI surrounding AI's implications for inequality, racial and gender bias, transparency and accountability to the public, human rights, safety, and more (Fjeld et al., 2020; Morley et al., 2019; Shneiderman, 2022). Proponents of this frame are likely more prone to advocate for stricter regulation and precaution in policymaking. Yet, there is uncertainty and debate about whether this frame is likely to influence policy in a meaningful way (Hickok, 2021; Morley et al., 2023; Taeihagh, 2021), given numerous barriers

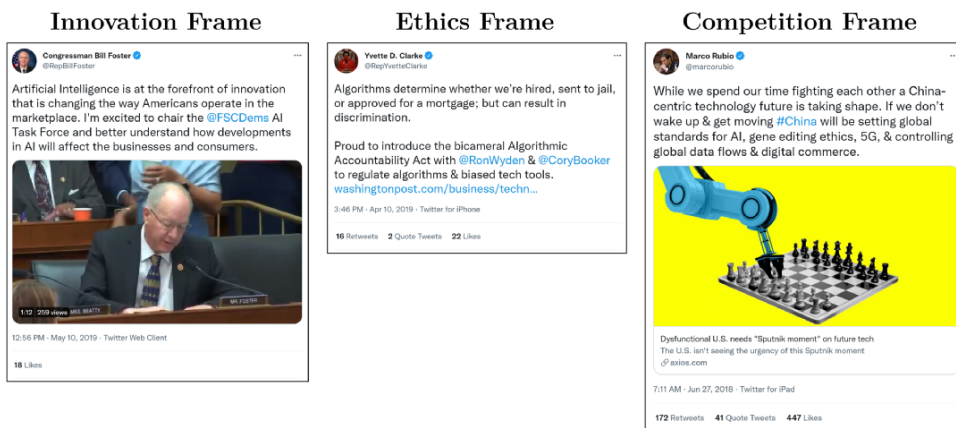


Figure 1. Examples of AI issue frames in US federal policy discourse.

to operationalizing AI ethics and concerns about private and public sector “ethics-washing” and “ethics shirking” (Bietti, 2020; de Laat, 2021).

The third frame is the “competition frame.”³ This frame—from the perspective of the US—largely highlights the “significant economic and national security threat to the United States” of losing the AI contest to China (Future of Defense Task Force, 2020, 5). There is a remarkable bipartisan consensus on this issue, with entities like the Future of Defense Task Force, the Office of Science and Technology Policy, the Congressional AI Caucus, and the National Security Commission on AI expressing support for the competition frame in their headline statements. Indeed, Stix and Maas (2021, 3) note that “public and global framings of AI in recent years have seemed to drift towards narratives of competition and ‘arms races’” and Imbrie et al. (2020, 8) argue that “the competition frame has diffused widely in public discourse and become shorthand for understanding the larger geopolitical context of investments in AI.”

Regarding the relative importance of these three frames, there are some reasons to believe that the ethics frame will be less influential than arguably more traditional frames invoking economic and security considerations. For instance, Gilardi et al. (2021) find that normative frames are less likely to predict policy adoption, while more “concrete” frames take their place increasingly over time. Moreover, economic and national security considerations typically rise to the top of public issue priorities (Jones et al., 2009). This, in part, results from—and reflects—the high levels of status and government access that actors focused on innovation and national security enjoy, compared to, for example, civil society groups.

Yet, the prevalence of the ethics frame in AI may be unprecedented in technology policy, surpassing ethical scrutiny applied to earlier emerging technologies like biotechnology and nuclear policy (Leung, 2020). Indeed, according to the European Commission in its landmark AI Act (European Commission, 2021, 8), ethics in AI is a “widespread and common approach, as evidenced by a plethora of ethical codes and principles developed by many private and public organisations..that AI development and use should be guided by certain essential value-oriented principles.” In sum, these considerations offer competing expectations with respect to the dominance of competing AI issue frames and their dynamics over time.

Ethical framing hypothesis

Compared to the innovation and competition frames, the ethics frame will find more purchase with the public, media, and policymakers, and is likely to become more dominant over time.

Traditional framing hypothesis

The innovation and competition frames will find more purchase with the public, media, and policymakers, and are likely to become more dominant over time.

Mutual interaction of public participation and framing dynamics

The ascendance of issue frames that capture social and ethical considerations would provide some evidence for transformation in technology governance, as would increased attention to public concerns. Yet it is the combination of the two that is arguably most indicative of a policy paradigm shift. It is not through sheer volume alone that policymakers are persuaded to change their policy preferences. Instead, frames may alter which actors have actual or perceived stake in, expertise relevant to, or even formal decision-making authority with respect to certain issues.

Particularly complex issues with low public salience admit to “board room politics” (Gormley, 1986), where citizens are disconnected from policy discussions while bureaucratic elites and regulatory capture characterize decision-making. In light of this, Lowi (1964) recognized that “one of the most important strategies in any controversial issue is to attempt to define it in redistributive terms in order to broaden the base of opposition or support.” Thus, issue framing can be used strategically to contain or expand issues and alter the actors that have influence, potentially facilitating or diminishing the influence of the public.

³ See Appendix B for a discussion of the selection of the three frames, including alternative approaches in the literature and a discussion of how frames are likely to emerge, disappear, merge, or shift over time. For instance, an “AI safety” frame has risen in popularity in recent years as evidenced by the launch of the US NIST AI Safety Institute and several international AI Safety Summits.

Table 1. Counts of issue frames used by public, policymakers, and media.

Dataset (2017–2019)	All AI messages	Ethics frame	Innovation frame	Competition frame
Public tweets	4,895,518	242,804	925,642	189,640
Congress tweets	1,477	413	417	187
NYT articles	5,773	3,921	5,170	3,957

In sum, if technology governance is shifting, we might expect that the increased risk aversion and precaution members of the public attach to emerging technologies is perceived as increasingly legitimate and worthy of serious consideration (de Saillie, 2015). Given wide calls for public participation in AI policy and evidence that the public is more skeptical of the use of AI than experts (O'Shaughnessy, et al., 2023), a key question is thus whether policymakers are indeed especially likely to listen to the public's social and ethical concerns regarding AI.

Special role of the public hypothesis

Issue attention by members of the public to the ethics frame will more strongly predict issue attention by policymakers.

Methodology

Data sources

There are three primary sources of data used in the study. To measure public attitudes towards AI, 4.9 million X (formerly Twitter) messages were collected that reference “#AI,” the standard hashtag for general AI interest. To measure policymaker discourse about AI, all messages on X (described here as tweets) sent by members of the 115th and 116th US Congresses were collected before extracting messages specifically about AI.⁴ Finally, while the relationship between the public and policymakers is of primary concern here, the media is incorporated primarily as a proxy of national mood or public opinion. Media attention to AI is captured by collecting all New York Times (NYT) articles that mention the phrase “artificial intelligence”⁵ using the Nexis Uni database. Additional discussion of the definition of “public,” choices surrounding collection of policymaker data, the role of the media, and the rationale for using social media data to study agenda-setting is available in Appendix A. All data cover the time period of January 2017 through December 2019, reflecting Congressional and executive priorities during a single presidential administration and key period in emerging AI policy. Notably, then, while this study provides insights on the earliest years of AI policy agenda-setting, it is circumscribed in its ability to assess the impact of public participation or issue framing beyond this time range, which entails that future research is needed to assess how these trends may evolve.

Identification of issue frames: text analysis

Quantitative content analysis was applied to each dataset to identify issue frame prevalence per actor (Boräng et al., 2014), primarily leveraging a dictionary approach.⁶ Ethics keywords were drawn from studies of AI conferences, policy documents, and media, while innovation and competition keywords came from media framing research focused on AI. Snowball searches and word embeddings were employed to find additional terms, and variations of terms were examined for relevance. A message or article that contains one or more of the frame-specific keywords is then indicated, non-exclusively, as portraying that respective frame. Additional details on dictionary creation and validation are available in Appendix C, full dictionaries for all datasets are available in Appendix D, and the sample sizes for each corpus and issue frame are displayed in Table 1.

Measuring agenda-setting influence: time series analysis

After constructing the datasets, the main analysis strategy used to study between-actor influence is time series analysis, particularly ARIMA and VAR modeling approaches. These methods offer

⁴ See Appendix A regarding the approach used to identify AI messages. The full dictionary for the 115th and 116th Congresses is available in Appendix D.

⁵ The choice of the different primary keywords is based on evaluation of the coverage of those terms, in light of differences in the length, context, and style of content.

⁶ Note that techniques like supervised classification could provide an alternative. See Appendix C for a discussion.

advantages such as modeling temporal order to support stronger causal claims, and are designed to deal with associated challenges such as autoregression in time series data and residuals as well as non-stationarity and conditional heteroscedasticity (Vliegenthart & Walgrave, 2008). ARIMA methods have been recommended for use in media and communication studies (Hollanders & Vliegenthart, 2008), and more recently in policy and political science scholarship, including in studies of agenda-setting using social media data specifically (Barberá et al., 2019), and in correlating media mentions with policy issue attention (Howlett, 1998).

Preparation for the ARIMA analysis includes examining the time series data and associated autocorrelation and partial autocorrelation functions. Analysis indicates evidence of trends but not seasonality, and data are aggregated weekly, in part because of sparsity in policymaker data. Next, analysis involves iterative search over possible stationary, non-seasonal ARIMA(p, d, q) specifications using maximum likelihood estimation and a small sample consistent Akaike Information Criterion (AIC) to balance predictive fit against model parsimony. Results suggest that models such as ARIMA(0, 1, 1) (simple exponential smoothing) and ARIMA(0, 1, 2) (damped trend exponential smoothing) may be appropriate.

To extend this analysis to bivariate regressions, it is necessary to model multiple time series simultaneously. The approach used here is ordinary least squares (OLS) regression using ARIMA to model errors, and including the predictor time series as an exogenous regressor. This strategy is permissible when the two series in a given regression are co-integrated, which is confirmed in all cases by the Phillips-Ouliaris test. Thus the basic specification is OLS regression with two actors' AI issue attention represented as either the independent or dependent variable, modeling a simple contemporaneous bivariate causal relationship within the course of a week. A benefit of this approach is that it allows for simple OLS-style interpretation. If the Public Agenda-Setting Hypothesis is correct, we would thus expect an increase in public messaging about AI to predict an increase in policymaker messaging.⁷ Appendix G provides further details about the preparation and justification for modeling choices.

While ARIMA is particularly suitable for bidirectional relationships where the order of causation is well known, VAR analysis allows for multidirectional influence through simultaneous estimation of multiple equations. While typically applied in macroeconomic analysis and forecasting, VAR has also been recommended and applied to analyze media and agenda-setting effects where causal relationships are less obvious and endogeneity is a concern (Gilardi et al., 2022; Liu et al., 2011). This is particularly important in policy scholarship given that policy theory often highlights multi-actor influence and complex feedback effects. However, while VAR analysis enables modeling of all three actors simultaneously, it emphasizes lagged rather than contemporaneous relationships, and interpretation can be more complex. I therefore follow the advice of Vliegenthart (2014) by considering the two methods as complementary approaches that may reveal different dynamics associated with issue framing and agenda-setting influence.

Importantly, coefficients cannot generally be interpreted straightforwardly for VAR. A few alternative methods are normally used. First, forecast error variance decomposition estimates indicate, for each time series, the amount of variation over time that can be attributed to its own lagged values as well as to the lagged values of the *other* endogenous variables (in this case, the other actors). Second, impulse response functions provide a graphical depiction of the movement of a response variable to a unit shock in the impulse variable. Appendix H provides further details on the modeling approach and decisions for the VAR Analysis, and Appendix I and J provide additional results.

Results

Issue frame prominence and trends

Figure 2 displays the proportion of all AI tweets (or articles) per actor that address each issue frame over the 2017–2019 time period (additional information is presented in Appendix F). As such, it depicts the *relative attention* to each issue frame over time out of all AI messages (which need not sum to 100%). A few patterns are clear. For the public, policymakers, and NYT, innovation appears to be the most prevalent frame, with a consistently high and stable level of attention from the beginning of the time period. In any given week, around 20–30% of public tweets, 20–50% of Congress tweets, and more than 90% of NYT articles which address AI discuss some aspect of innovation. The consistent attention across all actors to AI's implications for innovation is highly supportive of the Traditional Framing Hypothesis, though with respect to economic rather than geopolitical dynamics.

⁷ Further, we would expect the converse (policymakers influencing the public) to not be true, as this would raise concerns about endogeneity.

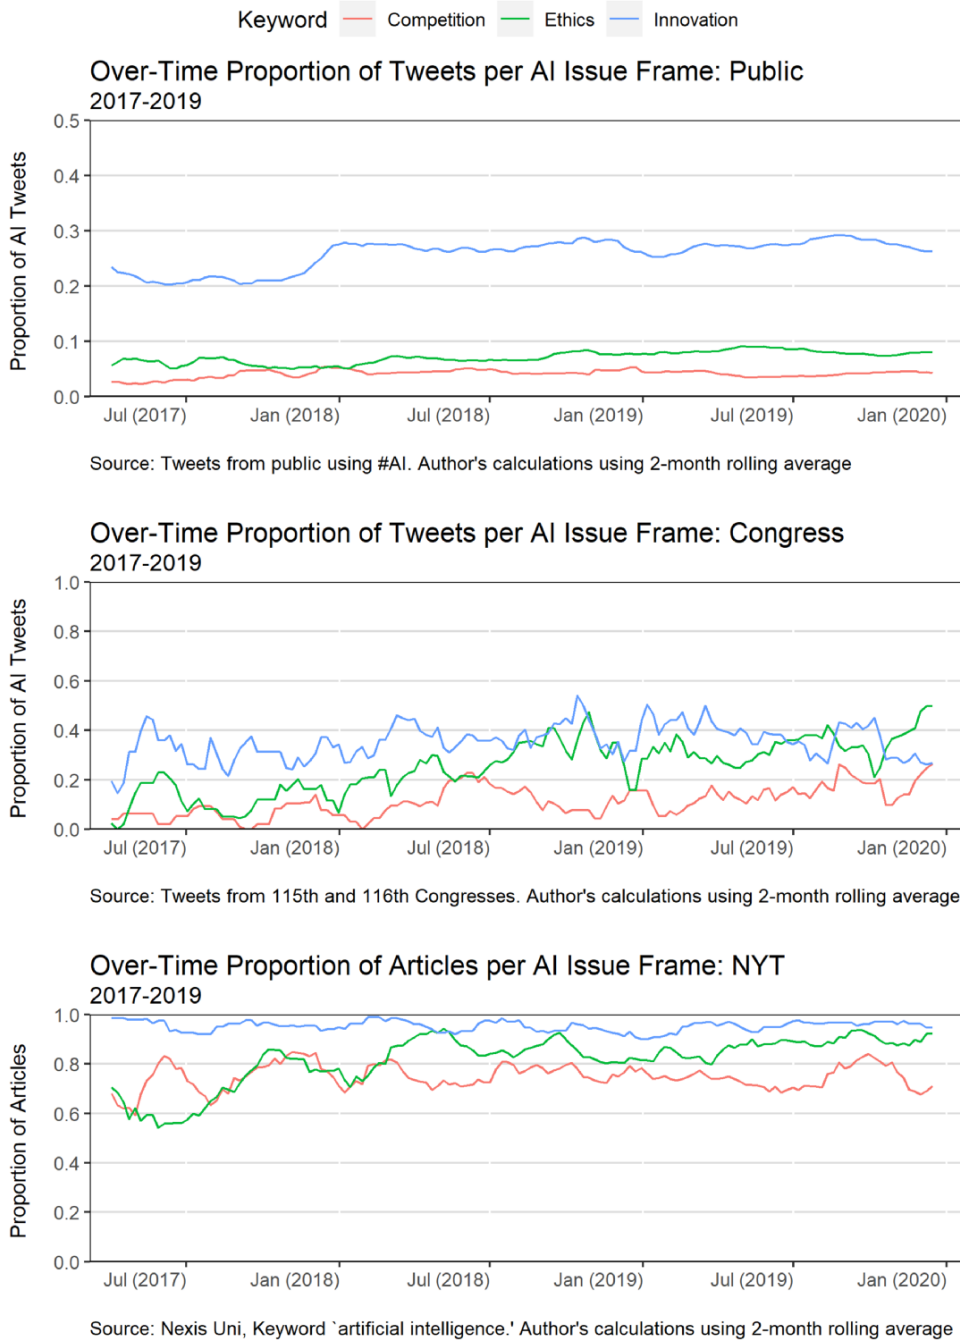


Figure 2. Issue frame prevalence by actor over time.

Note: The data cover January 2017 through December 2019. However, I use 2-month rolling averages to stabilize trends, which means the first data points depicted are from February 2017. Sources are (1) tweets from the public using #AI, (2) tweets from the 115th and 116th Congresses on AI, and (3) NYT article coverage of “artificial intelligence.”

Indeed, and in contrast, the competition frame receives the least attention overall. Yet, it appears that competition-related discourse is growing amongst members of Congress, increasing steadily from around 0% to around 25% of AI mentions between 2017 and 2019. A ramp-up of geopolitical discourse

about AI—including outside of the time window examined here—seems quite plausible given the centrality of US–China competition in post-2019 legislation such as the US Innovation and Competition Act. That the increase in attention to geopolitical dimensions seems particularly marked amongst members of Congress may also serve as a preliminary indication that the emerging policy agenda is forming somewhat independently from public influence (potentially contradicting the Public Agenda-Setting Hypothesis). The next subsections test this more formally.

The most striking dynamics, however, surround the ethics frame. Indeed, attention to the ethics frame appears to be increasing substantially over time, especially for the media and policymakers. Relative attention to the ethics frame steadily grows from around 70% of media discourse to more than 90% over the 3-year time period. In Congress, it increases from essentially 0% of AI discourse to approximately 50%, even surpassing attention to the innovation frame. This constitutes meaningful evidence that, notwithstanding the absence of attention to AI ethics during the early years of AI policy agenda-setting, ethics appears to have taken on a major role. Research extending into future time periods would be necessary to confirm these trends. Yet, attention to AI ethics does appear to be manifesting through legislation and other governance efforts like the Algorithmic Accountability Act, the AI Bill of Rights, and standards-related AI ethics efforts in the National Institute of Standards and Technology like the NIST AI Risk Management Framework and AI Safety Institute.

However, growth in Congressional attention to AI ethics does not seem to be tracking public attention, though there is modest growth in public attention to the ethics frame, from around 5–6% to 8%. The public's greater interest in AI innovation likely results from the particular composition of social media users who discuss AI broadly, who likely differ from the specialized mini-public (e.g., academia and civil society) focused on AI ethics. These results could signal an elite-public (and media-public) divide, implying that advocates of a new paradigm of technology governance may overstate or at least not closely reflect the current priorities of the public, particularly regarding social and ethical implications of AI. Alternatively, these trends could equally serve as a call for action to better inform and engage the public.

Overall then, the findings offer strong support for the Traditional Framing Hypothesis, particularly with respect to innovation rather than competition during this time period. Interestingly, the results also provide some support for the Ethical Framing Hypothesis given the striking growth in the ethics frame over time, which could have meaningful impacts on the AI policy agenda. However, there are early indications that the Special Role of the Public Hypothesis may not bear out.

ARIMA analysis

The ARIMA analysis involves modeling relationships between each pair of two actors separately for the three issue frames plus overall AI attention. Maximum likelihood estimation (MLE) is used to search through different ARIMA(p, d, q) specifications using a small sample consistent AIC to identify the best-performing non-seasonal model. Statistical and visual tests of goodness-of-fit for plausible models include examining autocorrelation and partial autocorrelations of residuals along with applying the Augmented Dickey Fuller test, Kwiatkowski–Phillips–Schmidt–Shin test, and Ljung-Box test, in order to confirm stationarity of the time series and independence of residuals. On balance given these diagnostics, a preferred model is selected, leaning towards simpler models over complex ones if they perform similarly, as it is preferable for no more than one of the p or q terms to be larger than 1 to avoid overfitting (Brandt & Williams, 2006).⁸

Table 2 displays the main ARIMA results from bivariate regressions of the three actors' issue and issue framing attention datasets. The first column indicates the actor (Y) potentially influenced by another leading actor (X). The second column indicates the selected ARIMA parameters based on information criteria and modeling diagnostics. The third column indicates the effect in terms of the *number of predicted additional tweets or articles per week* from Y that result from a one standard deviation increase in messages from X. The fourth column and fifth column assist with interpretation by presenting this

⁸ Note that this strategy means there are several different preferred p, d, q specifications for various bivariate relationships, when it could be argued that the same underlying processes should manifest across issue frames if agenda-setting dynamics are indeed part of a stable underlying system. As such, Appendix 1 presents key results using alternative ARIMA and fixed effects specifications as robustness checks. The fixed effects specifications differ in considering the impact of public tweets on individual Congressman behavior at the weekly level, but nevertheless lead to highly consistent results; impacts are even greater when applied to a subset of social media-engaged and AI-engaged legislators. This provides some reassurance in the robustness of the main models presented and in the underlying mechanisms studied here.

Table 2. ARIMA results: mutual influence of public, policymakers, and media.

X influence on Y	ARIMA model	Effect size	Baseline value	% Change	p-value
Public on Congress					
All	(0,1,1)	1.49	6.66	22.4%	0.003
Ethics	(0,1,1)	0.26	1.98	13.1%	0.149
Innovation	(0,1,1)	0.46	2.28	20.2%	0.015
Competition	(0,1,1)	0.03	0.81	3.7%	0.733
Congress on public					
All	(0,1,1)	753	31,322	2.4%	0.002
Ethics	(1,1,1)	60	2,205	2.8%	0.161
Innovation	(1,1,1)	136	8,046	1.7%	0.042
Competition	(1,1,1)	17	1,303	1.4%	0.694
NYT on Congress					
All	(0,1,1)	0.87	6.66	13.1%	0.024
Ethics	(1,1,3)	0.27	1.98	13.6%	0.079
Innovation	(0,1,1)	0.41	2.28	18.0%	0.012
Competition	(0,1,1)	0.03	0.81	3.7%	0.736
Congress on NYT					
All	(1,0,1)	1.57	18.56	8.5%	0.01
Ethics	(0,1,1)	0.94	15.3	6.1%	0.119
Innovation	(2,0,3)	1.59	17.71	9.0%	0.003
Competition	(0,1,1)	0.17	13.95	1.2%	0.74
NYT on public					
All	(0,1,2)	-30	31,322	-0.1%	0.896
Ethics	(1,1,1)	42	2,205	1.9%	0.269
Innovation	(1,1,1)	24	8,046	0.3%	0.709
Competition	(1,0,2)	-12	1,303	-0.9%	0.774
Public on NYT					
All	(0,1,1)	0.17	18.56	0.9%	0.828
Ethics	(0,1,1)	0.50	15.3	3.3%	0.384
Innovation	(3,0,0)	1.18	17.71	6.7%	0.091
Competition	(0,1,1)	-0.28	13.95	-2.0%	0.594

change in terms of a percentage change compared to the baseline average weekly value. For example, the results indicate that a one standard deviation increase in public tweeting behavior in a given week corresponds with a 22.4% increase over the baseline average of policymaker messaging that week. The final column indicates significance, though note that some statistically significant effects are too modest to be substantively important.

To assist with interpretation, [Figure 3](#) presents the main patterns of influence across the three actors and four issue frames. Larger magnitude effects are represented with bolder lines, and statistically significant effects at the 0.05 level are represented with solid rather than dashed lines. Overall, the results suggest that the public does influence policymakers, both overall and with respect to the innovation frame, leading to an approximate 20–22% increase in Congressional messaging over normal baseline attention. In contrast, even when the results are technically statistically significant, Congress exerts very minor influence on the public generally, on the order of 1–3%.⁹ Notably, attention by media to AI is also associated with increased attention by Congress on the order of 13–18%, indicating that this channel of influence is active, while attention by Congress leads to additional media attention (8–9%) as well. This confirms prior work recognizing the role of media in shaping and being shaped by policymaker attention ([Wolfe et al., 2013](#); [Yanovitzky, 2002](#)). To the extent that the media is treated as a reflection of public opinion, this provides further evidence of the Public Agenda-Setting Hypothesis.

Critically, both channels of public influence are present for all AI messages and for the innovation frame alone. In contrast—and across all actors—there are no significant patterns of influence for the ethics frame or for the competition frame at the 5% level. It is true, however, that the magnitude of

⁹ Indeed, note that for sets of bivariate relationships with less stable ARIMA parameters, generally insignificant results, and smaller effect sizes (e.g., Congress effects on public), this suggests that it is less likely that there is a stable and meaningful predictive relationship that can be captured.

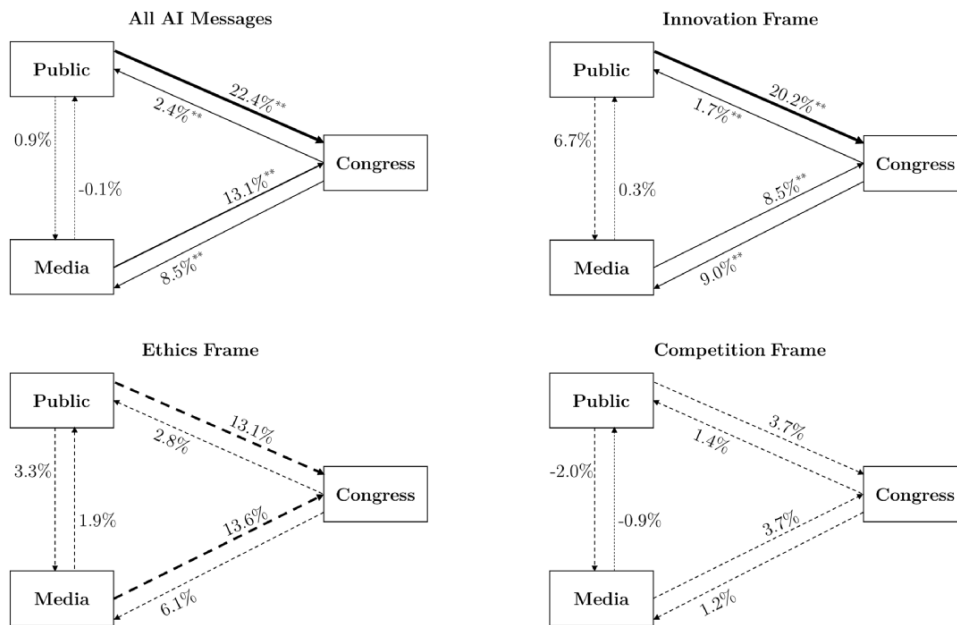


Figure 3. Results from ARIMA analysis: suggested bivariate relationships.

Note: Bivariate relationships per issue frame between public, policymakers, and media. Effect sizes listed correspond to percentage increase over weekly baseline resulting from one standard deviation increase in messaging from the influencer. Solid lines indicate significant results at the 5% level and line width corresponds to magnitude.

effects for the ethics frame is substantially larger and borderline significant at the 10% level in some cases, suggesting that the public could have some influence now or in the future. Meanwhile, the completely insignificant and marginal effects for the competition frame suggest the public has little role in shaping these dimensions of AI policy. It may be the case that policymakers are strictly looking to experts in geopolitical and military dimensions of AI on these issues. Overall then, while there is support for the Public Agenda-Setting Hypothesis with respect to the public's role generally, the *Special Role of the Public Hypothesis* with respect to the ethics frame is not supported. Instead, it appears that policymakers may listen to the public (and media) only when the economic and innovation-oriented dimensions of AI are emphasized.

While these results align to some degree with the descriptive trends, they should nevertheless be interpreted with caution. First, as they are unidirectional relationships that omit the influence of third party actors or other influences (such as focusing events), the results cannot safely be interpreted causally. Evidence of mutual influence in some cases could indeed indicate that both time series in a given relationship (e.g., public and Congress) are reflecting omitted external factors. Yet, it is plausible that there are not many additional external influences that would bring AI to the attention of policymakers or the public beyond what they learn from the media and one another. Further, because some of the relationships (again, for example, public and Congress) seem to strongly convey one-sided influence dynamics, it is less likely that confounders play a major role for these relationships. Another potential limitation is that the results depict contemporaneous or instantaneous influence within the course of a week, while other temporal modeling structures are possible.¹⁰ To address some of these concerns, especially those related to endogeneity, the next section presents results from the three-actor VAR analysis.

VAR analysis

Importantly, modeling with VAR allows for multidirectional influence and for more than two actors at a time. However, unlike with ARIMA, VAR measures lagged influence rather than instantaneous influence.

¹⁰ While a strong bidirectional relationship between two actors could indicate that the agenda-setting process is even shorter term and better modeled at the level of days, weekly structure is preferred here due to sparsity and to avoid having to specify an overly complex lag structure. Appendix G discusses temporal modeling choices in more detail.

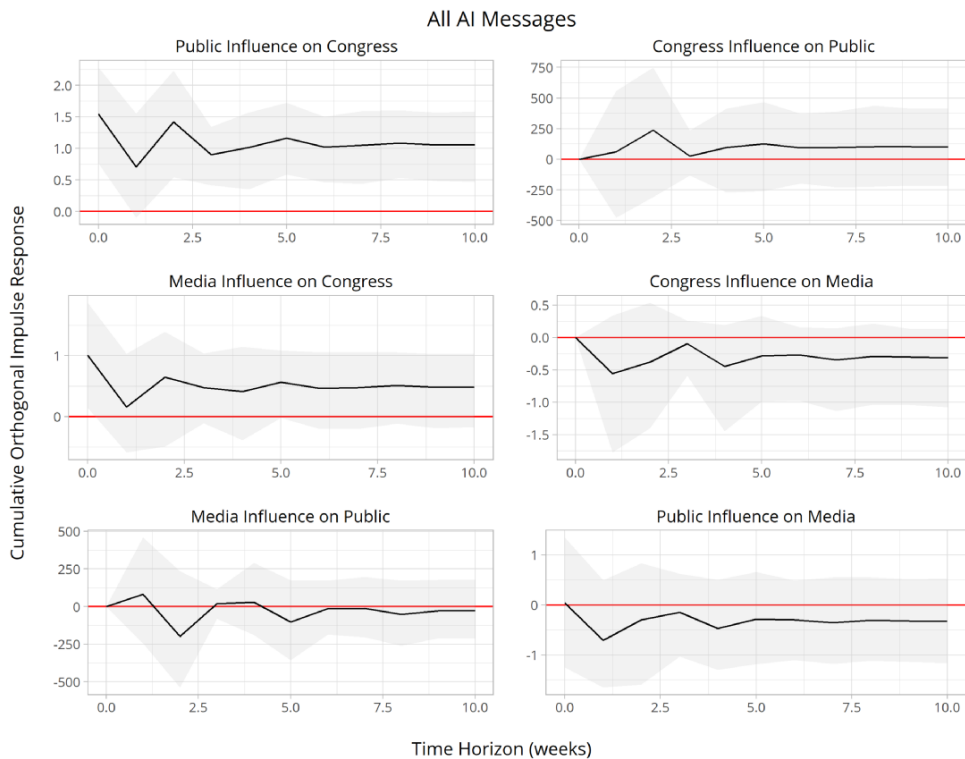


Figure 4. Mutual influence of public, policymakers, and media: all AI messages.

Note: Cumulative orthogonal impulse response functions indicating impact of a one standard deviation shock in the leading actor's attention to AI with 95% confidence intervals. Impacts measured in number of additional tweets or articles per week are along the Y axis; time in weeks is along the X axis.

A comparison of information criteria and prediction errors suggests two weekly lags as most appropriate. The VAR models again include all three actors (public, policymakers, and media) simultaneously, with separate analyses for the three issue frames plus AI attention overall.

Impulse response functions (additional results are available in [Appendix J](#), [Figures J1](#), [J2](#), and [J3](#)) are used to depict the impact of the *influenced* actor's attention to AI resulting from a one standard deviation shock in the *leading* actor's behavior. The results in [Figure 4](#) provide evidence that *public attention does influence policymakers, while Congress does not appear to influence the public*. The NYT also appears to have borderline influence on Congress, with a lesser magnitude compared to the public's influence, though bootstrap confidence intervals cross zero. Notably, results from the forecast error variance decomposition (available in [Appendix H](#)) also confirm these patterns, with the public explaining about 9–10% of the variance in Congressional AI attention overall, while Congress exerts no such influence on the public. Overall, there is further evidence of the Public Agenda-Setting Hypothesis, reiterating the prior ARIMA results.

However, the next key question is whether the public influences Congress with respect to *particular* issue frames. The next figure thus presents the impulse response analysis for the public specifically, for each issue frame and AI messages overall. The patterns in [Figure 5](#) roughly mirror the ARIMA results: *The public appears to have influence on Congress with respect to AI messages overall and for the innovation frame, but not for the ethics and competition frames.*¹¹ Complete results for each set of impulse response functions per issue frame across the three actors are presented in [Appendix H](#). In sum, there is little evidence for the Special Role of the Public Hypothesis.

¹¹ Note that while the magnitudes are small, e.g., 0.25 to 1.5 additional policymaker tweets about AI per week, this corresponds to around 10–20% increases over baseline Congressional messaging, as described in [Table 2](#).

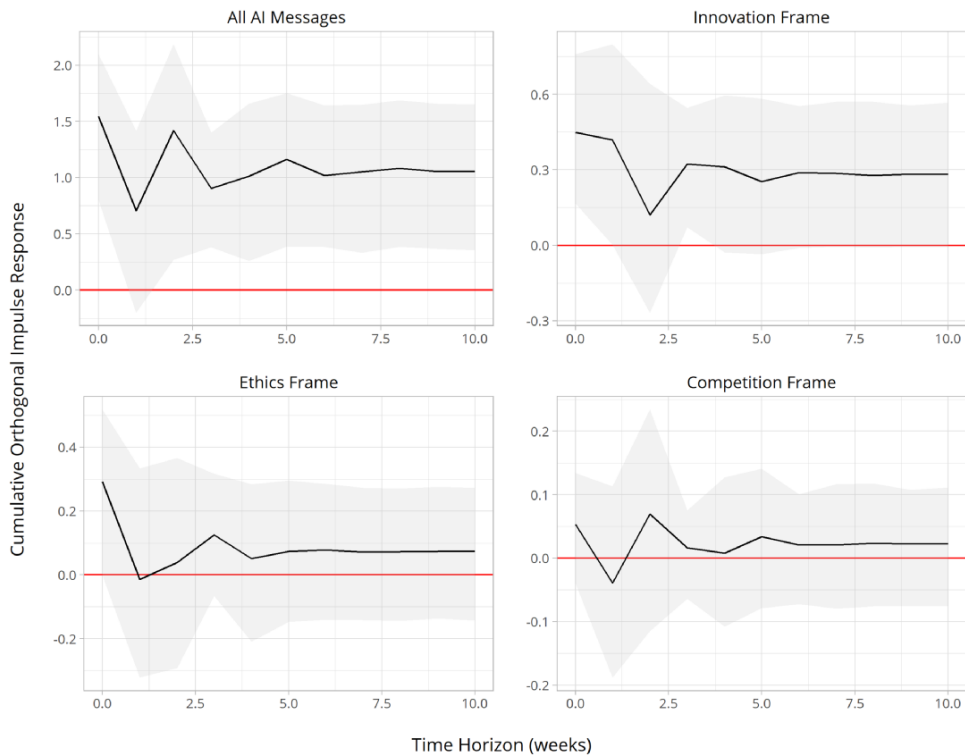


Figure 5. Public influence on policymakers per issue frame.

Note: Cumulative orthogonal impulse response functions indicating impact of a one standard deviation shock in public attention to AI with on policymaker attention with 95% confidence intervals. Impacts measured in number of additional tweets per week are along the Y axis; time in weeks is along the X axis.

Implications and conclusion

As policymakers increasingly wrestle with how to govern complex, high impact, strategic technologies, it is critical to understand both what kinds of issues will shape the policy agenda and which actors have influence in this agenda-setting process. In the context of AI policy, this paper considered whether the emerging AI agenda is developing along a traditional approach—emphasizing expert participation and strategic economic and geopolitical goals—or in line with a new paradigm surfacing social and ethical dimensions of technology and calling for increased public participation.

Findings here reveal a story that, while consistent, is also dynamic and mixed with respect to the role of issue framing contestation and public participation in AI policy. The descriptive results suggest that policymaker attention to the ethics frame for AI is indeed growing, even surpassing attention to geopolitical and economic dimensions of AI. These patterns may indeed be evidence that concerted efforts to govern technology with more proactive consideration of social and ethical dimensions are finding success. Further, across a variety of modeling approaches, results indicate that the public *does* lead policymaker attention to AI. This both contradicts and elaborates on prior work (Barberá et al., 2019) that did not find such a role for the general public across policy domains, while it did find a modest role for more attentive publics including aligned partisan publics, a finding this study cannot evaluate.

Increased public attention to AI is consistently associated with increased policymaker attention, with substantive effect sizes, while the converse is not true. Though less novel, the paper also finds that increased media attention, somewhat less consistently and to a lesser degree, is also associated with increased policymaker attention. The fact that public issue attention, directly or through media distillation, might influence the policy agenda even in the case of a highly complex and plausibly expert-dominated technology policy domain is striking.

Yet the prospects for meaningful public participation in AI governance also appear to be significantly circumscribed. Results demonstrate that the public influences policymaker attention to AI only when AI

is discussed in terms of its implications for economic growth and innovation. Meanwhile, and despite the activism of numerous actors in civil society and academia, the public does not appear to be an influence channel specifically with respect to social and ethical implications of AI. This could indicate that policymakers are engaged in a kind of confirmation bias, only listening to the messages they deem fit (Butler & Dynes, 2016). Alternatively, it could indicate that the public is relatively more invested in AI innovation than worried about AI's ethical risks, a finding reiterated in some public opinion research that shows stronger public support for adoption of AI than for its regulation (O'Shaughnessy et al., 2023). In either case, increased public participation and increased ethical consideration may not go hand-in-hand in a straightforward fashion, and imagining the public to serve as the "voice" of social and ethical consideration may be a simplification of a more complex and evolving role for the public. Additional strategies, like increased public education or facilitation of public involvement in policy via specialized fora, may be necessary to better understand and advance the public's role (Buhmann & Fieseler, 2023).

Importantly, there are several key limitations with respect to the study. First, the selection of issue frames and strategy for classification involve various subjective determinations. While studies aimed at identifying AI frames are broadly consistent with respect to general topics of importance (Imbrie et al., 2021; Ouchchy et al., 2020), other approaches to measuring issue framing could lead to different results. Second, though the 3-year time period under study is a critical one in light of the emergence of AI policy discourse at this time, it cannot reveal longer-term dynamics. For example, it may be the case that competition discourse has increased in the time period exceeding the scope of the study, that the public is gradually taking on a greater role as society becomes more familiarized with AI, that AI policy has been captured by industry, or that still other framings and policy actors have become ascendant. Research in future years and decades will be better positioned to determine if the dynamics identified here are stable or fleeting.

A third limitation surrounds the prospects for generalizing the findings here to other policy domains or settings. The results could indeed be signs of a paradigmatic shift in technology governance or agenda-setting generally, but this paper provides little direct evidence to that effect. Indeed, it is known that agenda-setting dynamics differ across policy domains, and AI could be a relatively unique issue given its technical complexity, widespread impact as a general purpose technology, and mixed benefits and risks. The US is also uniquely situated with respect to AI policy, as well as distinctive in terms of its agenda-setting institutions, policy goals, concerns, and so forth, meaning the exact findings may not translate to different regions and governance scales. Fourth, the time series models applied here do not explicitly take into account other factors like electoral context (Vliegenthart & Walgrave, 2008) or severity of problem indicators and focusing events (Liu et al., 2011). Given the short time period, subjectivity involved in time series modeling, and relatively sparse data in some cases (e.g., a small number of policymaker messages per week), future research will be needed to confirm or challenge the results here.

Overall, this study can benefit research in policy process theory, political communication, and AI policy with implications for the scholarly understanding of the roles of the public, media, and issue framing in the agenda-setting process. Through application of text-as-data methods and the creation of AI issue frame dictionaries, it also helps to concretize more preliminary and conceptual work concerned with AI issue frames and narratives, providing grounds for further study. Further, the application of time series methods to study not merely issue attention but also sub-issue frame attention by different policy actors represents a new approach to understanding the evolution of contested policy agendas. Finally, the paper helps to answer some concrete and pressing questions in technology and AI governance surrounding how AI is framed and whether the public has influence in the emerging agenda. While much remains to be known about the direction that technology governance will take, indications in the early years of AI policy suggest that a full transformation advancing the public's voice has not yet occurred.

Funding details

The author declares no funding.

Conflict of interest

None declared.

Appendix A

Data collection and rationale

I am able to identify Twitter accounts for 639 of 645 members of the Senate and House of Representatives—99% of those in office between 2017 and 2019—and collect tweets only during the specific terms when members were in office. As members of Congress often have multiple Twitter accounts (e.g., campaign accounts, personal accounts, official office accounts) (Siddique, 2019), I select for each person the account with the most followers, as this is the account with arguably the most importance for their messaging efforts. While 612 members of Congress tweeted approximately 1.1 million times over the period, only about half of them (291) tweeted about AI.

To identify AI-relevant messages, I investigate a set of 100 AI keywords (e.g., neural network, machine translation) used by Imbrie et al. (2020) for a similar task, along with additional keywords identified when investigating Congress's messages (e.g., autonomous vehicles, deepfakes). For each keyword, I investigate its frequency of occurrence, and retain keywords that appear at least 10 times in the corpus and for which at least 75% of the associated messages reflect a reference to AI, to my understanding.

A few additional questions are worth considering in light of the data collection strategy. First, to what extent does studying the sources above (namely social media) and actors allow for valid evaluations of issue (frame) priorities and agenda-setting influence? According to a growing body of work, these research questions are increasingly tractable. Twitter is now almost universally used by members of Congress, for purposes such as providing information, credit-claiming, and indicating responsiveness to the public, along with other goals (Golbeck et al., 2010; Shapiro et al., 2018). Research indicates that citizens also act strategically through these platforms to influence policy agendas (not limited to “shouting” at politicians) (Hemphill & Roback, 2014), and that policymakers in turn use social media to expand issue attention beyond elites as well (Fazekas et al., 2021). Perhaps surprisingly, nearly half of US politicians engage in conversation on Twitter; about one-quarter of this engagement is directly with private citizens (Tromble, 2018).

A related question is whether issue attention on social media is reflective of issue attention more broadly. Regarding issue congruence across sources, evidence indicates that issue attention by policymakers on Twitter is indeed substantially similar to their issue priorities in other formats such as on Facebook, via press releases, and in parliament (Casas & Morar, 2019; Peeters et al., 2021). These priorities are also congruent with issue coverage in the NYT specifically, suggesting that the NYT is a useful source for studying elite-media issue attention relationships (Shapiro & Hemphill, 2017). Thus, there are good reasons to think that issue attention, issue congruence, and responsiveness between the public, policymakers, and media can be plausibly evaluated using the identified datasets, even though the platforms, contexts, and rationales for each actor's messaging efforts are somewhat distinct. In sum, there is growing evidence of the importance of social media in agenda-setting dynamics.¹² Social media may serve as a direct channel in influence efforts, and—given issue congruence across platforms—it may also serve as a meaningful proxy of public and policymaker opinion. This paper considers both channels to be viable. However, much remains unknown about the role of social media in agenda-setting, especially given the novelty of the platforms and the still-evolving political communication dynamics surrounding them. This paper helps to advance this literature.

A distinct question important in the context of the study is what is meant by “public.” Public may refer to any actors outside of government (e.g., industry bodies and academic experts who reply to requests for public comment); it may refer to society or citizens generally; or it may refer to special subsets of the public (Hallahan, 2000). For example, various taxonomies distinguish between “strong” and “weak” publics, between “inattentive” and “attentive” or “engaged” publics (Fraser, 2021; Stoker, 2014), and scholars dispute whether there is such a thing as a general public or rather multiple mini-publics. Notably, many definitions focus on the shared stake in and potentially action surrounding particular social or policy issues (Barnes et al., 2003), meaning that public is often defined as explicitly policy-engaged. Along these lines, this paper considers the relevant public to be individuals on social media who have some heightened personal stake or interest in AI and are more likely to engage politically on those issues.¹³

¹² See Barberá et al. (2019) and Gilardi et al. (2022) for further discussion.

¹³ Notably, this Twitter audience is likely to skew male, be younger, more educated, live in urban areas, and is potentially more politically liberal than the general population (Barberá & Rivero, 2015; Mellon & Prosser, 2017; Mislove et al., 2011). Yet, such a public is not dissimilar from individuals more likely to contact political representatives, vote, attend hearings,

Finally, an important note is that the NYT is often used in media studies as a proxy for media coverage broadly as it is considered the “paper of record” in the US. It covers a wide range of topics and is widely read by the public and political elites (Ringel, 2021; Weaver & Bimber, 2008). In the agenda-setting literature, the media plays a complex and important role (McCombs & Shaw, 1972). It may serve to synthesize societal debate and reflect overall public opinion (Hopkins et al., 2017; Ripberger, 2011), or may indeed *construct* public opinion (McCombs, 2004; Zaller, 1992). Further, media outlets are not only responsive to the needs of their customers, but are also responsive to political elites, both by reporting on their activity and providing crucial information to them (Callaghan & Schnell, 2001; Shoemaker & Reese, 1991; Van Aelst & Walgrave, 2016). Given the importance of public opinion and national mood in the agenda-setting literature, I consider the media here primarily as a proxy or alternative source that policymakers look to understand public opinion. Indeed, other research on AI policy has observed the prominent role of the media in presenting issue frames (Imbrie et al., 2021) and has considered media coverage in outlets like the NYT as a proxy for public opinion (Fast & Horvitz, 2017; Ouchchy et al., 2020).

Appendix B Selection of issue frames

The choice of frames is, to some extent, subjective and constrained by which audiences and time period are considered (Fast & Horvitz, 2017). For example, the Center for Security and Emerging Technology considers four frames including a “World without Work” frame and a “Killer Robots” frame (Imbrie et al., 2021) and the Department of Homeland Security (Department of Homeland Security, Office of Cyber and Infrastructure Analysis, United States, 2017) considers seven benefit- and threat-focused frames including “Threat to Humanity” and “Fueling the Surveillance Machine” frames.

The choice of frames for this analysis is based on review of scholarly and policy literature along with expert opinion and a desire for parsimony. Approaches like discourse analysis and topic modeling can be fruitfully applied to study coverage of AI in media and other sources by those who wish to identify, combine, and parse frames in different ways (Cave et al., 2018). Yet, Neuman et al. (1992) identify human impact, economics, and conflict as three common issue frames used in the news media; arguably, then, the three frames pursued in this study have universal relevance.

It is also important to note that frames may evolve over time as particular frames emerge, disappear, become subsumed or merged into other frames, or simply change in meaning. For instance, in the 2020s, there has been some ascendance of an “AI safety” frame, which contains elements related to an “AI ethics” frame but also incorporates considerations related to existential risk or the “Threat to Humanity” frame as identified by the Department of Homeland Security in its analysis.

Appendix C Identification of issue frames

Within each of these three datasets, quantitative content analysis methods are used to identify the prevalence of issue frames (Boräng et al., 2014).¹⁴ To identify the ethics frame, I draw on keywords from three studies used to comprehensively identify AI ethics concepts in AI conferences and journals (Prates et al., 2018), in public, private, and NGO sector AI policy and ethics documents (Schiff et al. 2021), and in media sources (Fjeld et al., 2020; Zhang et al., 2021). For the competition and innovation frames, I rely on keywords used by Imbrie et al. (2020) and Imbrie et al. (2021) in their extensive study of AI frames in media coverage. For all three frames, I engage in snowball searches of the tweets/articles through an iterative process to identify and evaluate other potentially applicable terms. Finally, I employ word embeddings via the word2vec family of algorithms (Mikolov et al., 2013) to identify other candidate terms.

In constructing the three dictionaries, for each possible term, I examine possible variations of the associated stem of that term to determine relevance. For example, the terms “moral” and “develop” capture ethics and innovation concepts effectively, while the terms “morale” and “developers” do not. I deem a keyword worthy of inclusion if at least 75% of the usages of it in a random sample of messages accurately reflect the relevant concept in AI, based on my substantive understanding of the field.

etc. (Vaccari et al., 2015). That is, a public so defined is a highly relevant audience for considering political accountability, responsiveness, and agenda-setting influence.

¹⁴ I begin with standard cleaning and pre-processing steps for text analysis, including removing duplicate tweets/articles, text segmentation, tokenization and lemmatization, removal of punctuation and numbers, case conversion, and collapsing of some n-word terms to individual tokens.

I perform this analysis for both the Twitter and NYT datasets, as the nature of term usage depends on the length and context of the medium (Grimmer & Stewart, 2013), and find the three dictionaries perform adequately overall compared to similar classification efforts (Ghosh & Loustaunau, 2021).¹⁵ A message (tweet or article) that contains one or more of the frame-specific keywords is then indicated, non-exclusively, as portraying that respective frame. The full dictionaries are available in Appendix D and the sample sizes for each corpus and issue frame are displayed in Table 1.

Appendix D

AI and issue frame dictionaries

This section presents the dictionaries used for extracting AI-related messages and the three issue frames studied in this paper. While there were some minor differences in the initial keywords and validation efforts with respect to the different data sources examined in this study, the resulting dictionaries were substantially similar, so only one dictionary is presented for AI overall and for the three issue frames. Of note, the dictionaries were adjusted throughout the research process as classification improved; however, the main results were not sensitive to these minor changes. Nevertheless, differences in data sources, initial search criteria, length of texts, time period, etc., imply that it is important for researchers to verify and potentially adjust these AI dictionaries when applying them to another context or data source (Grimmer & Stewart, 2013).

In addition to the dictionaries themselves, the high-level categories listed here can also be used or repurposed to extract content related to specific sub-dimensions of AI discourse that may be of interest to researchers, such as “discrimination” or “military.”

Dictionary for identifying AI discourse:

- [AI]: ai*, a.i.*, artificial* intellig*, artificialintelligen*, artificial-intell*
- [automation]: automate, automated, automation, autonomous syst*, autonomous tech*
- [ML]: machine learn*, deep learn*, reinforce* learn*, supervised learn*, unsupervised learn*
- [algorithm]: algorithm*
- [other_techniques]: recommend* system*, speech recognition, computer vision,
- pattern recognition, natural language, machine translation, recommendation system*, image process*, information retrieval, speech synthesis, handwriting recognition, object recognition
- [AVs]: autonomous veh*, self-driv*, self driv*, autonomous car*, AVs, autonomous navigation, AV START, AV test*, AV tech*
- [facial_recognition]: facial recog*
- [deepfakes]: deepfake*, deep fake*
- [weapons]: autonomous weapon*

Dictionary for identifying AI’s economic and innovative dimensions—“innovation frame”:

¹⁵ Note that this initial classification strategy only identifies term-specific true positives and false positives, but does not provide a holistic sense of false negatives or other classification criteria. As such, I hand code a subset of documents and compare accuracy against the dictionary-based classifications. I report results in Appendix E. Note that techniques like supervised classification could provide an alternative (Wilkerson & Casas, 2017). However, a study using a wide variety of state-of-the-art ML methods, including various doc2vec and self-attention based classifiers to specifically study AI frames in media documents, led to a maximum F1 classification score for the positive class of 0.81 for the best classifier, with precision and recall of 0.76 and 0.87 for the positive class (Ghosh & Loustaunau, 2021). For this reason and because these techniques have not been applied to study shorter Twitter content for AI issue framing, I stick with the simpler dictionary approach, which Gentzkow et al. (2019, 24) expect to “remain the optimal choice in many settings.”

- [innovation]: innovat*, grow, growth, patent*, invent*, invest*, billion*
- [economy]: economy, economies, economic
- [transformation]: transform*, chang*, revolution*, disrupt*, reinvent*, disrupt*, reinvent*, shake
- [acceleration]: accelerat*, drive, drives, driven, enhanc*, boost*, fueled, fueling, enable, enabler*, enables, enabling, optimize, optimizing, optimizes, augment*
- [business]: business, businesses, start-up*, startup*, enterpris*
- [opportunity]: embrac*, opportunity, opportunities, ambitions, potential
- [productivity]: productiv*, efficien*, edge
- [pace]: pace, faster, momentum, rise

Dictionary for identifying AI's social and ethical dimensions—"ethics frame":

- [ethics]: ethic*, moral*, morality, evil
- [rights]: rights, freedom*
- [values]: human values, societal values, democratic values
- [responsibility]: responsibil*, responsibleAI
- [control]: human control
- [fairness]: fairness*, fairer*, unfair*
- [discrimination]: discrim*, non-discrim*, anti-discrim*, antidiscrim*
- [transparency]: transparenc*, explainab*, interpretab*
- [safety]: safe*
- [accountability]: accountab*
- [privacy]: privac*
- [bias]: bias*, unbiased, racis*, sexism*, homophob*, transphob*
- [equality]: equality*, unequal*, unequal*
- [equity]: equita*
- [humanity]: humanity, humane, human-cent*
- [trust]: trust, trustworth*, trusted, untrust*, fear*, afraid*, threat*
- [misuse]: misus*, abus*
- [manipulation]: manip*

- [harm]: harm, harmed, harmful, harming, harms
- [consent]: consent*
- [deception]: deceiv*, decept*, fooled
- [misinformation]: misinform*, propagand*, democracy
- [displacement]: displac*, unemploy*
- [diversity]: diverse, diversity
- [society]: societ*
- [sustainability]: sustainable, sustainability, environmental*

Dictionary for identifying AI's strategic, geopolitical dimensions—"competition frame":

- [competition]: competit*, compete, rivals, rivalry
- [position]: ahead*, left behind*, fall* behind, advantag*, rank, ranks
- [supremacy]: dominance, dominate, dominating, domination, supremacy
- [china_russia]: china*, chinese*, russia*
- [arms_race]: arms race, artificial intelligence race, AI race, global race, #AI race, sputnik
- [war]: war, warfare, battlefield*, weapon*, lethal*
- [military]: national security, militar*, army*, navy*, naval, air force*, marines, armed forces

Appendix E Classification results

To evaluate the quality of the AI issue frame dictionaries, I perform a manual classification on 1,000 randomly selected messages from the public and separately on 1,000 randomly selected messages from Congress. The results in [Table E1](#) report the precision (or positive predictive value) and recall (or sensitivity) for each dataset and issue frame. I emphasize these metrics and the overall F1 score, given this study's focus on predicting the positive class, i.e., identifying the issue frames themselves.

Overall, the classification leads to adequate to strong performance for both social media datasets with F1 scores ranging from 0.81 to 0.96.¹⁶ The innovation frame is captured most accurately across datasets, with slightly lower performance for the ethics frame in the public dataset and for the competition frame in the policymaker dataset. Of note, the dictionaries appear to perform slightly better on the policymaker messages, perhaps because these messages are, on average, more carefully crafted as well as focused on a smaller subset of AI issues.

¹⁶ Note that I do not perform classification on the NYT articles given that they are of secondary importance in this study and because of the more pressing need to evaluate the robustness of the dictionaries in the context of short Twitter messages. For readers interested in classifying news media messages with respect to AI issue frames, [Ghosh and Loustaunau \(2021\)](#) provide some promising approaches.

Table E1. Classification results for issue frames: public and Congress datasets.

Issue frame	Precision	Recall	F1 score
Public AI messages			
Ethics	0.96	0.70	0.81
Innovation	0.98	0.95	0.96
Competition	0.86	0.91	0.89
Policymaker AI messages			
Ethics	0.98	0.87	0.92
Innovation	0.93	0.93	0.93
Competition	0.98	0.79	0.87

Appendix F

Issue frame timelines and descriptive statistics

Figure F1 depicts the raw correlations between the three actors' time series data across the three issue frames plus AI messages overall. Across all three actors, there is modest correlation with respect to "all AI messages" and for the ethics and innovation frames, but little correlation with respect to the competition frame. This provides preliminary evidence that AI issue attention is generally correlated across actors, though these raw correlations do not account for spurious trends and autocorrelation of each time series with itself.

Beyond the question of whether single frames "prevail" during a specific time period, frames may also merge over time or become subsumed into one another. Policymakers might attempt to advance both ethics and competition simultaneously, such as by arguing that ethical AI is a strategy to promote competitive advantage (Minkkinen et al., 2023).¹⁷ Additionally, policymakers may employ the ethics and innovation frames simultaneously, such as by arguing that ethical AI is needed to foster trust and support consumer adoption.¹⁸ Gilardi et al. (2021) find evidence that frames tend to become more complex over time—effectively by merging together. These changes may reflect the development of a stable consensus and dominant "policy image" (Baumgartner & Jones, 1991), at least in the near-term.

Along these lines, Figure F2 depicts the prevalence of issue frames over time for each actor while additionally incorporating hybrid or joint frames, e.g., *messages that discuss both ethics and innovation*. Overall, the vast majority of messages contain just a single issue frame, and there is little evidence of growth for any hybrid frame, particularly amongst the public and Congress. Because NYT articles are longer than tweets, however, they are more likely to reflect both individual and hybrid frames given the dictionary classification approach utilized in the study. However even here, the hybrid frames seem to merely track the growth of individual frames rather than represent a new framing consensus.

Appendix G

Notes on ARIMA analysis approach

The ARIMA(p, d, q) models considered in this analysis take the general form:

$$y'_t = a_0 + \sum_{i=1}^p a_i * y_{t-i} + \sum_{i=1}^q b_i * \epsilon_{t-i} + \epsilon_t \quad (1)$$

where y'_t is a variable that has been differenced from itself one or more times to remove trends or other statistically undesirable properties that might lead to spurious correlations, (e.g., $d = 1$ corresponds to $y' = y_t - y_{t-1}$); p refers to the number of autoregressive terms, or lagged values of y ; q refers to the number of moving average terms, or lagged prediction errors estimated from previous time periods; and ϵ refers to the unexplained random residual. Selecting the appropriate ARIMA(p, d, q) specification

¹⁷ For example, the European Commission's Ethical Guidelines for Trustworthy AI (European Commission, 2019, 5) state that ethical AI confers a "competitive advantage" for individual firms and EU countries: "A trustworthy approach is key to enabling 'responsible competitiveness', by providing the foundation upon which all those affected by AI systems can trust that their design, development and use are lawful, ethical and robust."

¹⁸ For example, according to US Executive Order on Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government (White House, 2020, 1): "The ongoing adoption and acceptance of AI will depend significantly on public trust. Agencies must therefore design, develop, acquire, and use AI in a manner that fosters public trust and confidence while protecting privacy, civil rights, civil liberties, and American values."

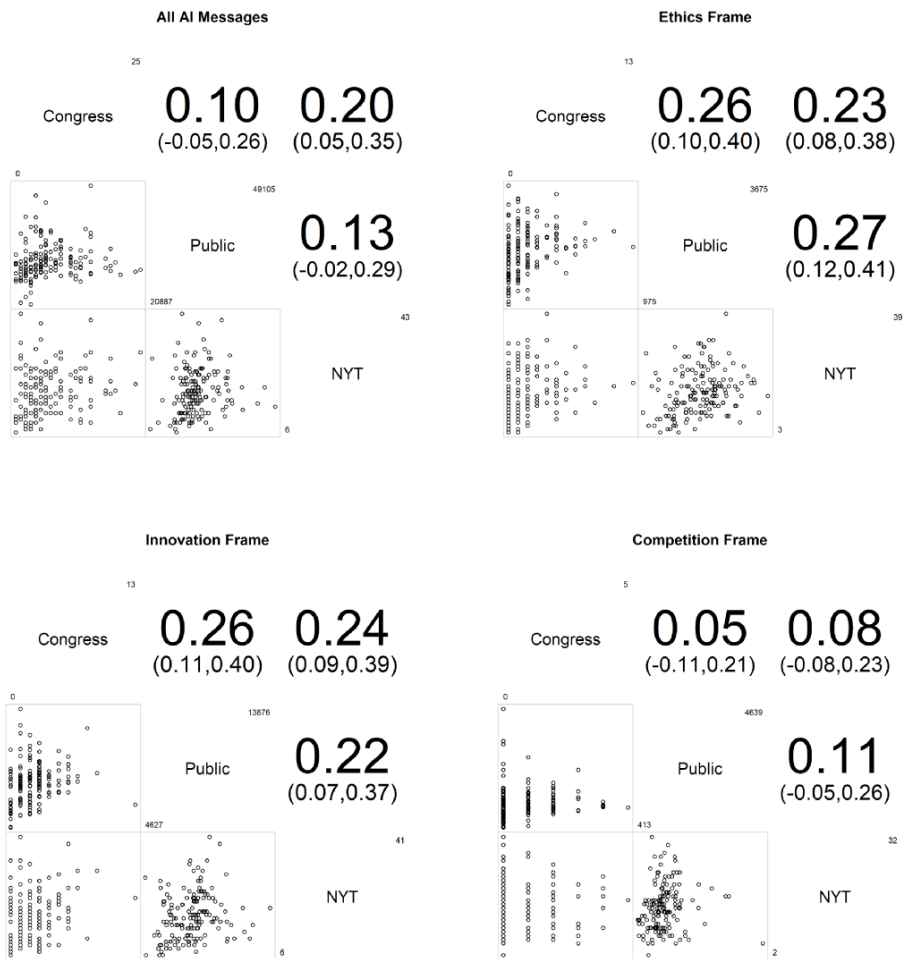


Figure F1. AI issue frame correlations across public, policymakers, and media.

Note: Depicts raw correlations between the 12 time series datasets: three issue frames plus overall AI messages across the three actors during 2017–2019. There are small correlations for AI messages overall and for the innovation and ethics frames.

typically requires understanding the autocorrelation and partial autocorrelation functions for a given time series, and examining information criteria such as the Akaike Information Criterion (AIC) and Schwarz Bayesian Criterion (BIC) to identify which model produces stronger goodness of fit, balanced against ensuring parsimony in the number of terms to avoid overfitting. I prefer the use of the AIC or AICc (an AIC with a small sample correction) (Bozdogan, 1987) as underfitting is a more significant problem for small sample sizes (Dziak et al., 2012; Vaida & Blanchard, 2005).²⁰ A resulting model should have residuals that are uncorrelated and stationary, with constant error variance over time.

There are a total of 12 time series involved in this analysis (three datasets × mentions for each of three issue frames plus total number of mentions). Not all time series may have the same properties, which complicates model identification and choice of analytical method. However, analysis of seasonality using trend and seasonal decompositions suggests that there are not seasonal patterns or cycles, so I use non-seasonal analyses throughout. Visual evidence and tests (the Augmented Dickey-Fuller Test or ADF test and Kwiatkowski–Phillips–Schmidt–Shin or KPSS test) also suggest at least some time series (e.g., the ethics frame) are characterized by over-time trends, which implies that one level of differencing

²⁰ Nevertheless, using alternative information criteria like the BIC does not lead to substantively different results across the main analyses.

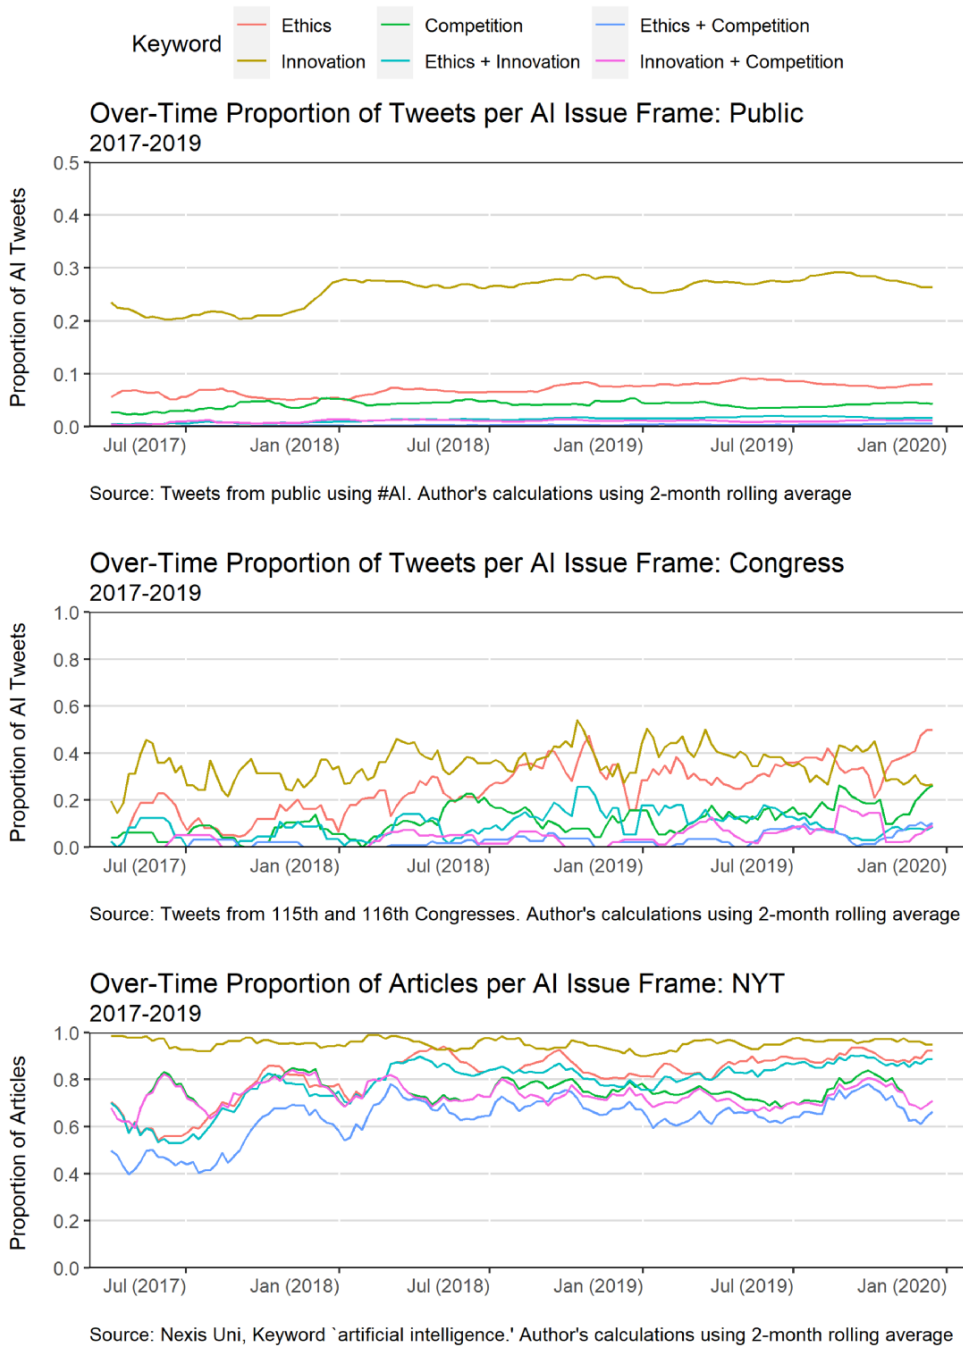


Figure F2. Issue frame prevalence by actor with joint frames.

Note: Single and joint issue frames from January 2017 through December 2019. I use 2-month rolling averages to stabilize trends, which means the first data points depicted are from February 2017. Sources are (1) tweets from the public using #AI, (2) tweets from the 115th and 116th Congresses on AI, and (3) NYT article coverage of "artificial intelligence."

may be appropriate to remove those trends and render the time series stationary, a prerequisite for time series modeling.

I aggregate the time series data to the level of week for three reasons. First, some times series data in my dataset, particularly policymaker messaging, is relatively sparse at the daily level. Second, messaging behavior indicates visible within-week relationships, such as a tendency to tweet the same day each week or more regularly on Fridays, and aggregation to the week level should smooth out these cycles. Third, theory and prior literature suggest that responsiveness to messages may occur over days-to-weeks rather than more quickly (hours) or slowly (months) (Gonzales, 2018). For example, news media and policymakers may only start to message about new legislation, hearings, or focusing events over several days rather than immediately.

Another question is whether the models should consider the leading actors to influence other actors simultaneously (i.e., within the same week) or with a lag (e.g., only a week or two later). Models performed based on 1-, 2-, or 3-week lags exclusive of contemporaneous influence exhibit decreasing and insignificant effects.

Additionally, multi-model selection identifies contemporaneous influence models as the most plausible, and does not recommend including prior week lags at all. Tests of instantaneous and Granger causality using multi-actor VAR modeling also suggest that contemporaneous rather than past influence is important, especially for AI messages overall and for the innovation frame (with no evidence of instantaneous influence for the ethics and competition frames). And as stated, I expect attention to AI issue frames to manifest over fairly short time frames. Thus, I adopt a contemporaneous influence modeling approach using ARIMA for the primary specifications, while I consider lagged influence more directly in the VAR analysis.

Using a small-sample adjusted AICc as the primary selection criteria, I search iteratively over possible stationary, non-seasonal models using maximum likelihood estimation (MLE) to maximize fit. Results suggest that models such as ARIMA(0, 1, 1) (simple exponential smoothing) and ARIMA(0, 1, 2) (damped trend exponential smoothing) may be appropriate (McKenzie & Gardner, 2010) though the preferred model varies by time series. This strategy generally results in stationary models and residuals, i.e., independent errors without unit roots in the moving average or autoregressive terms. This provides some insight about the nature of the time series under study.

However, to extend this analysis to bivariate regressions, it is necessary to model multiple time series simultaneously. I therefore vary the standard ARIMA approach applied to a single time series and depicted in the equation above. The approach used here is ordinary least squares (OLS) regression using ARIMA to model errors, and including the predictor time series as an exogenous regressor. Thus the basic specification is a simple bivariate OLS with two of the three actors represented as independent or dependent variables:

$$y_t = \beta_0 + \beta_1 x_t + \epsilon_t \quad (2)$$

but with ϵ_t modeled in the style of Appendix G, and with the x_t and y_t terms potentially subject to differencing to remove trends.

A benefit of this approach is that it allows for simple OLS-style interpretation, while avoiding problems associated with autoregressive errors. The AICc is used to identify the final model for each relevant bivariate regression of two actors, and differencing is applied to both time series variables in almost all of these models.

Importantly, it is allowable to apply this approach even to nonstationary time series if some linear combination of those series is stationary, i.e., if the series are “cointegrated.” Thus, I do not engage in extra pre-cleaning steps beyond differencing. The Phillips-Ouliaris (PO) test confirms that the series in the study are cointegrated in all cases. Also note that, while there is information about the past incorporated in each model, the modeling strategy implies that the effects of x on y are instantaneous or contemporaneous.

Appendix H

Notes on VAR analysis approach

In contrast to ARIMA, a VAR model allows for lagged values of the key dependent variable as well as lagged values of separate exogenous variables. The standard or reduced form of a VAR(1) representation (with a single lag) is:

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} + \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} + \begin{bmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{bmatrix}$$

such that $y_{1t} = a_{11}y_{1t-1} + a_{12}y_{2t-1} + \epsilon_{1t}$ and $y_{2t} = a_{21}y_{1t-1} + a_{22}y_{2t-1} + \epsilon_{2t}$. Thus, all key time series variables (e.g., public, policymaker, and media messages) are included in both the left- and right-hand sides of the respective equations, in a symmetric fashion.

Diagnostic tests (ADF, KPSS, and PO) provide mixed results regarding whether the relevant time series are stationary. For this reason, differencing (of all series) may be appropriate, though some research cautions that differencing will remove long-term dynamics (Brandt & Williams, 2006) and suggest using Vector Error Correction Models (VECM) instead. As long-term dynamics are not the central focus of this analysis, I stick with VAR analysis of differenced series. Next, I use the AIC to determine the appropriate lag order, p for a VAR(p) process, as AIC is again a preferable information criteria for smaller samples. A total of two to three (week) lags is deemed appropriate depending on whether differencing is applied, and seems sufficiently parsimonious. The results from VAR models are consistent with no autocorrelation or heteroskedasticity in the residuals, as desired.

Figure H1 presents the forecast error variance decomposition based on all AI messages, indicating what proportion of the variance of a given actor's behavior over time is explained by lagged values of its own time series versus that of other actors. The results indicate that AI attention in the media and public is substantially autonomous. For both the public and NYT, only 1–2% of AI messaging behavior is explained by other actors. This is consistent with prior scholarship demonstrating a high degree of attentional inertia (Liu et al., 2011; Wood & Peake, 1998). In contrast, approximately 9–10% of policymaker AI messaging is explained by prior public messaging, with some influence (4–5%) from the media as well. Thus, when accounting for multidirectional influence and all three actors, the findings align with main ARIMA results and the Public Agenda-Setting Hypothesis.

Appendix I

Alternative specifications: fixed effects and ARIMA models

As an additional set of robustness checks, I alternatively run fixed effects models to assess the impact of public AI issue attention on Congressman AI issue attention overall and per issue frame. This

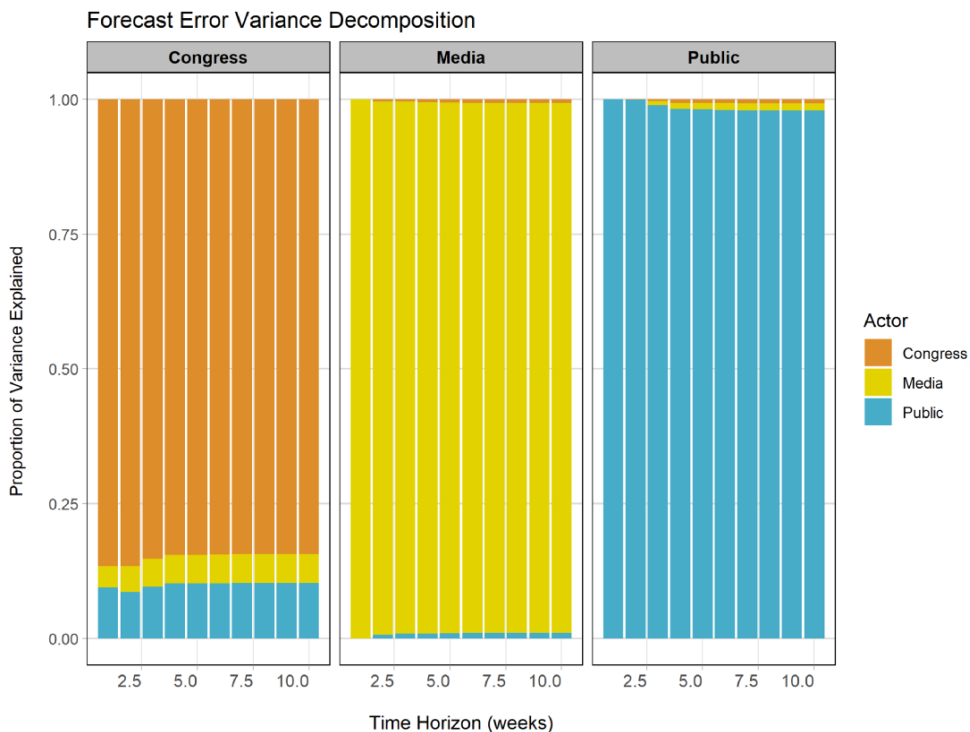


Figure H1. Forecast error variance decomposition: mutual influence of actors.

Table 11. Public impact on policymaker attention: fixed effects specifications.

	Congress AI issue attention			
	All AI messages (1)	Ethics (2)	Innovation (3)	Competition (4)
Public: AI messages	0.002*** _(0.001)			
Public: ethics frame		0.0002 _(0.0002)		
Public: innovation frame			0.001*** _(0.003)	
Public: competition frame				0.00003 _(0.0001)
Week trend	0.0001*** _(0.00001)	0.00005*** _(0.00001)	0.00002*** _(0.00001)	0.00002*** _(0.00000)
Congressmember fixed effects	Yes	Yes	Yes	Yes
Observations	100,323	100,323	100,323	100,323
R ²	0.053	0.030	0.034	0.024

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

approach is quite distinct from the ARIMA and VAR models. Rather than aggregating all Congressmembers' messaging behavior into a single weekly count (representing policymaker attention as a whole), I consider impacts on individual *Congress members per week*. Specifically, I include Congressmember fixed effects in my model specification, and cluster standard errors by week as the "treatment" of public AI issue attention is delivered at the level of weeks. Finally, I include a simple linear time trend to account for the order of weeks.

The results in [Table 11](#) indicate that a one standard deviation increase in public attention to AI is associated with a statistically significant average increase of 0.0017 tweets per Congressmember per week. Multiplied by the members of Congress in the dataset, this corresponds to 1.10 additional tweets per week across all of Congress. Similarly, increased public messaging about the innovation frame is associated with about 0.45 additional weekly tweets by Congress related to AI innovation, quite similar to the effect sizes found with the ARIMA models. Yet, also as with the ARIMA results, the public is only influential with respect to all AI messages and the innovation frame. In contrast, public messaging about ethics during a given week is associated with a very modest 0.099 increased tweets by Congress, while messaging about competition leads to a minuscule 0.021 increased tweets by Congress, and both of these latter associations are highly insignificant ($p = 0.47$ and $p = 0.75$, respectively).

The fact that this quite distinct modeling structure leads to similar findings with respect to public influence provides further confidence about the robustness of the ARIMA results. However, I consider two additional robustness checks using unit fixed effects based on Congressmember-week AI attention. First, I subset the sample to members of Congress who are particularly engaged in social media. In particular, I restrict the sample to members of Congress who are above the median in terms of total social media messaging over the study's time period. If public-policymaker engagement over social media indeed constitutes a valid channel for agenda-setting influence, we should expect these social media-engaged legislators should be particularly influenced by increased public attention.

The results in [Table 12](#) appear to confirm this expectation. Multiplying by the number of engaged legislators, increased public messaging is associated with 6.23 additional weekly tweets about AI across all social media-engaged members of Congress ($p < 0.01$), and 4.18 additional tweets about AI innovation ($p < 0.01$), but has minuscule and insignificant effects for the ethics frame (0.63 additional tweets, $p = 0.72$) and competition frame (0.22 additional tweets, $p = 0.87$).

Finally, I restrict the sample to only AI-engaged legislators, identified as legislators who tweeted about AI 10 or more times during the 3-year period. These legislators are amongst the most active on AI policy, including members of the House and Senate Congressional AI Caucuses who are often co-sponsors of AI-related legislation and members of committees that discuss AI in Congressional hearings and reports. Again, to the extent that the public does indeed influence policymaker issue attention to AI, we should expect that these members of Congress would be especially attentive and responsive.

The results in [Table 13](#) provide further confirmation that AI-engaged policymakers are particularly responsive to the public. Across all of Congress, increased public messaging is associated with 5.45 additional weekly tweets about AI for all AI-engaged members of Congress ($p < 0.01$) and 6.05 additional tweets about AI innovation ($p < 0.01$), but still has relatively far smaller and insignificant effects for the

Table 12. Public impact on policymaker attention: social media-engaged policymakers.

	Congress AI issue attention: active tweeters			
	All AI messages (1)	Ethics (2)	Innovation (3)	Competition (4)
Public: AI messages	0.020 ^{***} (0.006)			
Public: ethics frame		0.002 _(0.006)		
Public: innovation frame			0.014 ^{**} (0.006)	
Public: competition frame				0.001 _(0.005)
Week trend	0.001 ^{***} (0.0002)	0.001 ^{***} (0.0002)	0.001 ^{***} (0.0001)	0.001 ^{***} (0.0001)
Congressmember fixed effects	Yes	Yes	Yes	Yes
Observations	48,042	48,042	48,042	48,042
R ²	0.053	0.031	0.034	0.025

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.Note: Subset to $n = 306$ legislators who are actively engaged on Twitter, or legislators who tweeted above the median amount of times amongst legislators over the 3-year time period.**Table 13.** Public impact on policymaker attention: AI-engaged policymakers.

	Congress AI issue attention: AI-engaged legislators			
	All AI messages (1)	Ethics (2)	Innovation (3)	Competition (4)
Public: AI messages	0.130 ^{***} (0.046)			
Public: ethics frame		0.017 _(0.038)		
Public: innovation frame			0.144 ^{***} (0.052)	
Public: competition frame				0.015 _(0.031)
Week trend	0.006 ^{***} (0.001)	0.008 ^{***} (0.001)	0.002 ^{**} (0.001)	0.005 ^{***} (0.001)
Congressmember fixed effects	Yes	Yes	Yes	Yes
Observations	6,594	6,594	6,594	6,594
R ²	0.037	0.029	0.029	0.023

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.Note: Subset to $n = 42$ "AI-engaged" legislators, or legislators who tweeted about AI at least 10 times over the 3-year time period.

ethics frame (0.73 additional tweets, $p = 0.64$) and competition frame (0.62 additional tweets, $p = 0.63$). Overall, these additional specifications not only help to confirm the key findings emerging from the time series models; they also provide some additional evidence that public issue attention is especially influential for policymakers who are engaged in AI policy, and suggest that public influence on the policy agenda may indeed be operating directly through social media as a channel.

One concern with respect to the use of time series models to assess multiple bivariate relationships is that the identified "preferred" (p, d, q) model differs in some cases. Yet, to the extent that agenda-setting influence behavior is part of a stable underlying data generating process or processes, we might expect that the same kinds of temporal dynamics should be in play across actors and issue frames. Therefore, Table 14 reproduces the main results in Table 2, but using a consistent time series specification for all actors and issue frames. In particular, I use an ARIMA(0, 1, 2) or damped trend exponential smoothing approach to model the errors.

The results are highly stable here and with respect to other minor adjustments in the ARIMA modeling approach not displayed here. There are no major changes to the significance of the particular between-actor influence relationships, and effect sizes are largely similar as well, providing some confidence in the robustness of the main paper results.

Table 14. ARIMA(0, 1, 2) results: mutual influence of public, policymakers, and media.

X influence on Y	ARIMA model	Effect size	Baseline value	% Change	p-Value
Public on Congress					
All	(0,1,2)	1.48	6.66	22.2%	0.004
Ethics	(0,1,2)	0.26	1.98	13.1%	0.152
Innovation	(0,1,2)	0.47	2.28	20.6%	0.015
Competition	(0,1,2)	0.03	0.81	3.7%	0.728
Congress on Public					
All	(0,1,2)	753	31,322	2.4%	0.002
Ethics	(0,1,2)	60	2,205	2.7%	0.163
Innovation	(0,1,2)	147	8,046	1.8%	0.036
Competition	(0,1,2)	19	1,303	1.5%	0.679
NYT on Congress					
All	(0,1,2)	0.9	6.66	13.5%	0.019
Ethics	(0,1,2)	0.26	1.98	13.1%	0.119
Innovation	(0,1,2)	0.43	2.28	18.9%	0.009
Competition	(0,1,2)	0.03	0.81	3.7%	0.733
Congress on NYT					
All	(0,1,2)	1.41	18.56	7.6%	0.024
Ethics	(0,1,2)	0.94	15.3	6.1%	0.116
Innovation	(0,1,2)	1.4	17.71	7.9%	0.012
Competition	(0,1,2)	0.18	13.95	1.3%	0.73
NYT on Public					
All	(0,1,2)	-30.01	31,322	-0.1%	0.896
Ethics	(0,1,2)	40	2,205	1.8%	0.291
Innovation	(0,1,2)	15	8,046	0.2%	0.813
Competition	(0,1,2)	-16.23	1,303	-1.2%	0.695
Public on NYT					
All	(0,1,2)	0.3	18.56	1.6%	0.702
Ethics	(0,1,2)	0.45	15.3	2.9%	0.441
Innovation	(0,1,2)	0.52	17.71	2.9%	0.563
Competition	(0,1,2)	-0.17	13.95	-1.0%	0.759

Note: ARIMA Model refers to the preferred (p, d, q) specification. Effect size refers to the number of additional AI messages by the influenced actor (Y) in a given week resulting from a one standard deviation increase of messaging from the influencer (X) during that week. Baseline value refers to the average number of AI messages from the influenced actor (Y) per week. % Change refers to the increase in AI messaging over the baseline resulting from a one standard deviation increase of messaging from the influencer during that week. Statistical significant results at the 5% level are highlighted. Note that some digits are trimmed for legibility.

Appendix J Additional VAR results

The additional figures presented here depict orthogonal cumulative impulse response functions across the three actors and three main issue frames. Results complement main Figure 4 and confirm that there are few relationships between two actors reflecting issue attention influence. That is, in addition to the public's influence with respect to AI issue attention *overall* on Congress, only Figure J2 depicting the innovation frame shows clear evidence of meaningful influence. This reiterates the study's general findings related to the Public Agenda-Setting Hypothesis and lack of evidence for the Special Role of the Public Hypothesis.

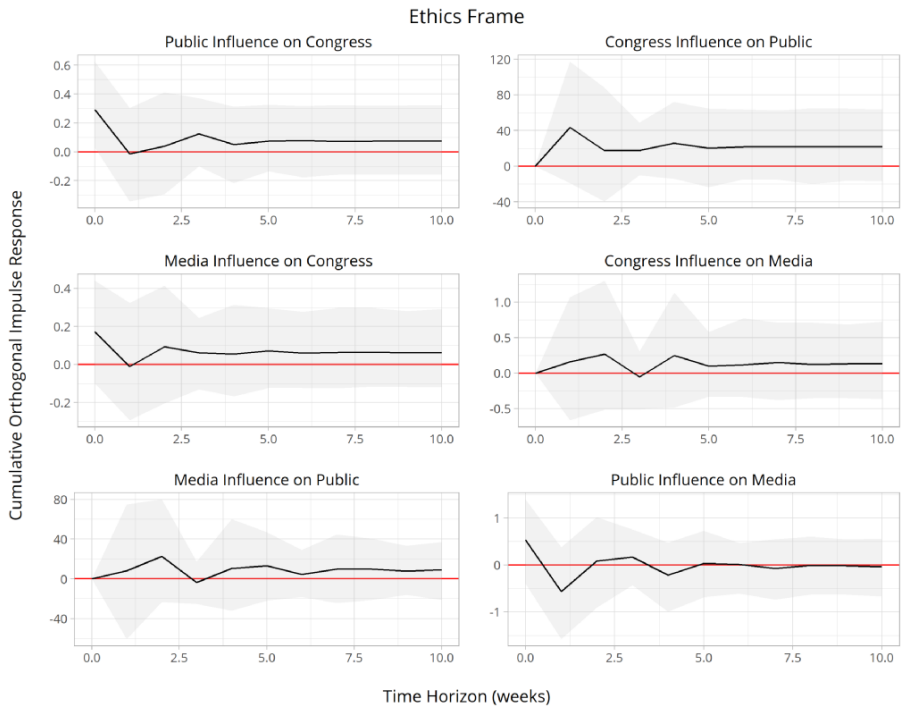


Figure J1. Cumulative impulse response functions: ethics frame.

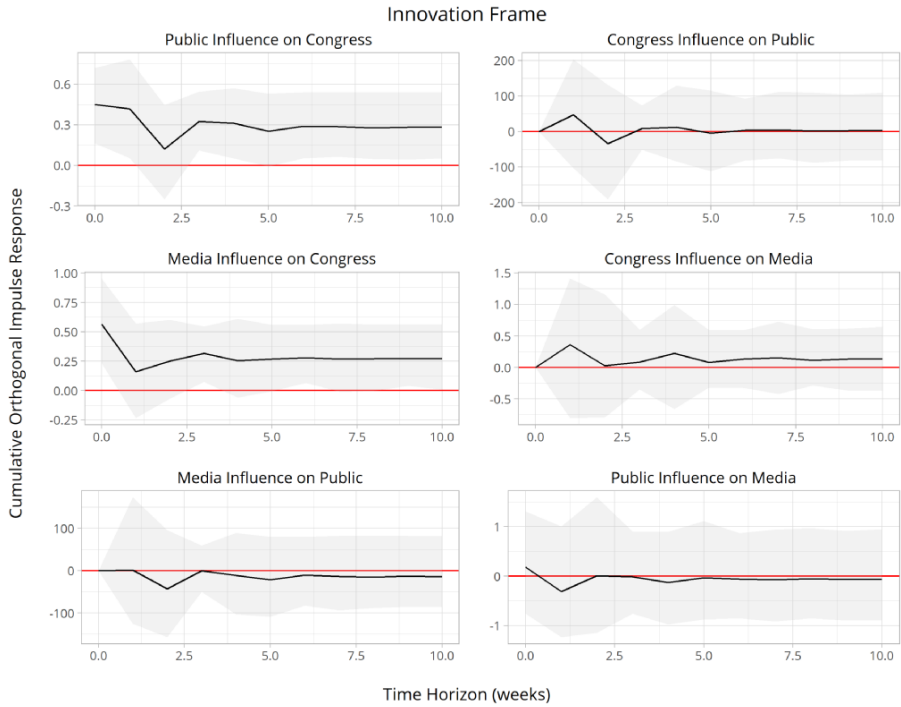


Figure J2. Cumulative impulse response functions: innovation frame.

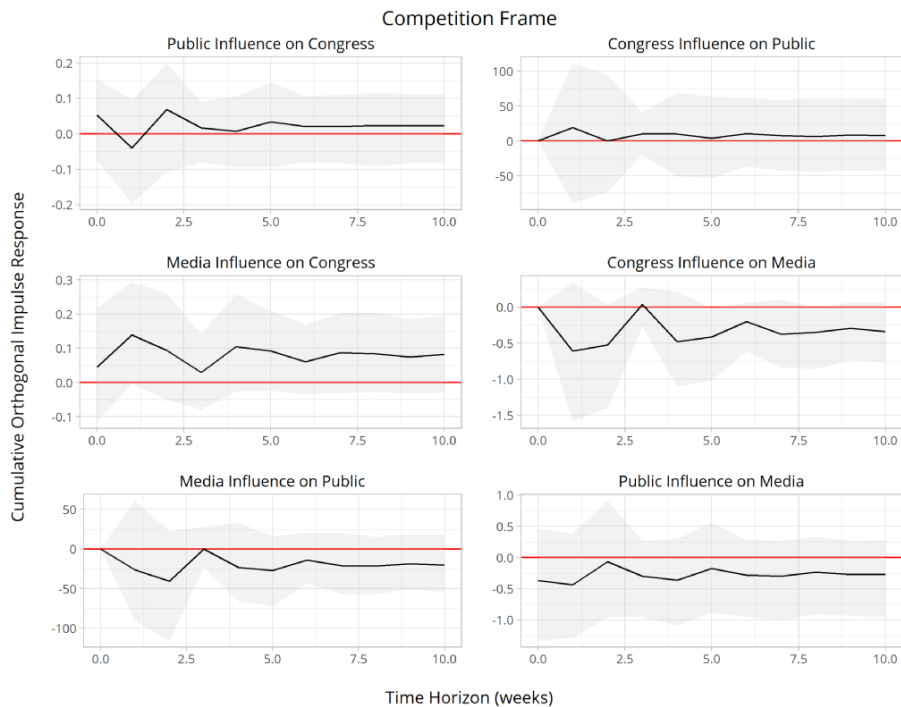


Figure J3. Cumulative impulse response functions: competition frame.

References

- Abelson, J., & Gauvin, F.-P. (2006). *Assessing the impacts of public participation: Concepts, evidence and policy implications*. Technical report, Centre for Health Economics and Policy Analysis (CHEPA), Hamilton.
- Barberá, P., Casas, A., Nagler, J., Egan, P. J., Bonneau, R., Jost, J. T., & Tucker, J. A. (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. <https://doi.org/10.1017/S0003055419000352>
- Barberá, P., & Rivero, G. (2015). Understanding the political representativeness of twitter users. *Social Science Computer Review*, 33(6), 712–729. <https://doi.org/10.1177/0894439314558836>
- Barnes, M., Newman, J., Knops, A., & Sullivan, H. (2003). Constituting ‘the public’ in public participation. *Public Administration*, 81(2), 379–399. <https://doi.org/10.1111/1467-9299.00352>
- Baumgartner, F. R., & Jones, B. D. (1991). Agenda dynamics and policy subsystems. *The Journal of Politics*, 53(4), 1044–1074. <https://doi.org/10.2307/2131866>
- Baumgartner, F. R., & Jones, B. D. (1993). *Agendas and instability in American politics*. In *American politics and political economy series* (1st edn.). University of Chicago Press.
- Bietti, E. (2020). From ethics washing to ethics bashing: A view on tech ethics from within moral philosophy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT**20*, pp. 210–219, Barcelona, Spain. Association for Computing Machinery.
- Boräng, F., Eising, R., Klüver, H., Mahoney, C., Naurin, D., Rasch, D., & Rozbicka, P. (2014). Identifying frames: A comparison of research methods. *Interest Groups & Advocacy*, 3(2), 188–201. <https://doi.org/10.1057/iga.2014.12>
- Bozdogan, H. (1987). Model selection and Akaike’s Information Criterion (AIC): The general theory and its analytical extensions. *Psychometrika*, 52(3), 345–370. <https://doi.org/10.1007/BF02294361>
- Brandt, P. T., & Williams, J. T. (2006). *Multiple time series models*. SAGE Publications.
- Buhmann, A., & Fieseler, C. (2023). Deep learning meets deep democracy: Deliberative governance and responsible innovation in artificial intelligence. *Business Ethics Quarterly* 33, 146–179. <https://doi.org/10.1017/beq.2021.42>

- Butler, D. M., & Dynes, A. M. (2016). How politicians discount the opinions of constituents with whom they disagree. *American Journal of Political Science*, 60(4), 975–989. <https://doi.org/10.1111/ajps.12206>
- Callaghan, K., & Schnell, F. (2001). Assessing the democratic debate: How the news media frame elite policy discourse. *Political Communication*, 18(2), 183–213. <https://doi.org/10.1080/105846001750322970>
- Casas, A., & Morar, D. (2019). *Different channel, same strategy? Filling empirical gaps in Congress literature*. SSRN Scholarly Paper ID 3395307, Social Science Research, Network.
- Cave, S., Craig, C., Dihal, K., Dillon, S., Montgomery, J., Singler, B., & Taylor, L. (2018). *Portrayals and perceptions of AI and why they matter*. Technical report, The Royal Society.
- Chong, D., & Druckman, J. N. (2007). Framing theory. *Annual Review of Political Science*, 10(1), 103–126. <https://doi.org/10.1146/annurev.polisci.10.072805.103054>
- Crockett, K., Colyer, E., & Latham, A. (2021). The ethical landscape of data and artificial intelligence: Citizen perspectives. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)* Orlando, FL, pp. 1–9.
- de Laat, P. B. (2021). Companies committed to responsible AI: From principles towards implementation and regulation? *Philosophy and Technology*, 34(4), 1135–1193. <https://doi.org/10.1007/s13347-021-00474-3>
- Department of Homeland Security, Office of Cyber and Infrastructure Analysis, United States. (2017). *Narrative analysis research paper: Artificial intelligence risk to critical infrastructure*. Technical report. Department of Homeland Security, Office of Cyber and Infrastructure Analysis.
- de Saille, S. (2015). Innovating innovation policy: The emergence of ‘Responsible Research and Innovation’. *Journal of Responsible Innovation*, 2(2), 152–168. <https://doi.org/10.1080/23299460.2015.1045280>
- Diercks, G., Larsen, H., & Steward, F. (2019). Transformative innovation policy: Addressing variety in an emerging policy paradigm. *Research Policy*, 48(4), 880–894. <https://doi.org/10.1016/j.respol.2018.10.028>
- Dryzek, J. S., & Pickering, J. (2017). Deliberation as a catalyst for reflexive environmental governance. *Ecological Economics*, 131, 353–360. <https://doi.org/10.1016/j.ecolecon.2016.09.011>
- Dziak, J. J., Coffman, D. L., Lanza, S. T., & Li, R. (2012). *Sensitivity and specificity of information criteria*. Technical Report 12-119, The Methodology Center, The Pennsylvania State University.
- Edler, J., & Fagerberg, J. (2017). Innovation policy: What, why, and how. *Oxford Review of Economic Policy*, 33(1), 2–23. <https://doi.org/10.1093/oxrep/grx001>
- European Commission. (2019). *Ethical guidelines for trustworthy AI*. Technical report, European Commission, High-Level Expert Group on Artificial Intelligence (AI HLEG).
- European Commission. (2021). *Proposal for a regulation on a European approach for artificial intelligence*. Technical Report COM(2021) 206 final, European Commission.
- Fast, E., & Horvitz, E. (2017). Long-term trends in the public perception of artificial intelligence. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI’17*, pp. 963–969, San Francisco, California, USA. AAAI Press.
- Fazekas, Z., Popa, S. A., Schmitt, H., Barberá, P., & Theocharis, Y. (2021). Elite-public interaction on twitter: EU issue expansion in the campaign. *European Journal of Political Research*, 60(2), 376–396. <https://doi.org/10.1111/1475-6765.12402>
- Fischer, F. (2000). *Citizens, experts, and the environment: The politics of local knowledge*. Duke University Press.
- Fischer, S.-C., Leung, J., Anderljung, M., O’Keefe, C., Torges, S., Khan, S. M., Garfinkel, B., & Dafoe, A. (2021). *AI policy levers: A review of the U.S. government’s tools to shape AI research, development, and deployment*. Technical report, Centre for the Governance of AI, Future of Humanity Institute, University of Oxford.
- Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). *Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI*. SSRN Scholarly Paper ID 3518482, Berkman Klein Center for Internet & Society.
- Fraser, N. (2021). Rethinking the public sphere: A contribution to the critique of actually existing democracy. In M. Mitrašinović & V. Mehta (Eds.), *Public space reader* (pp. 34–41). Routledge.
- Future of Defense Task Force. (2020). *Future of Defense Task Force report 2020*. Technical report, House Armed Services Committee.
- Gamson, W. A., & Modigliani, A. (1989). Media discourse and public opinion on nuclear power: A constructionist approach. *American Journal of Sociology*, 95(1), 1–37. <https://doi.org/10.1086/229213>
- Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3), 535–574. <https://doi.org/10.1257/jel.20181020>
- Ghosh, S., & Loustaunau, P. (2021). RHeFrameDetect: A text classification system for automatic detection of rhetorical frames in AI from open sources. *arXiv:2112.14933 [cs]*. <https://arxiv.org/abs/2112.14933>
- Gilardi, F., Gessler, T., Kubli, M., & Müller, S. (2022). Social media and political agenda setting. *Political Communication* 39(1), 39–60. <https://doi.org/10.1080/10584609.2021.1910390>

- Gilardi, F., Shipan, C. R., & Wüest, B. (2021). Policy diffusion: The issue-definition stage. *American Journal of Political Science*, 65(1), 21–35. <https://doi.org/10.1111/ajps.12521>
- Goddard, S. E., & Krebs, R. R. (2015). Rhetoric, legitimation, and grand strategy. *Security Studies*, 24(1), 5–36. <https://doi.org/10.1080/09636412.2014.1001198>
- Golbeck, J., Grimes, J. M., & Rogers, A. (2010). Twitter use by the U.S. Congress. *Journal of the American Society for Information Science and Technology*, 61(8), 1612–1621. <https://doi.org/10.1002/asi.21344>
- Gonzales, A. (2018). Do more with Twitter data.
- Gormley, W. T. (1986). Regulatory issue networks in a federal system. *Polity*, 18(4), 595–620. <https://doi.org/10.2307/3234884>
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297. <https://doi.org/10.1093/pan/mps028>
- Hallahan, K. (2000). Inactive publics: The forgotten publics in public relations. *Public Relations Review*, 26(4), 499–515. [https://doi.org/10.1016/S0363-8111\(00\)00061-8](https://doi.org/10.1016/S0363-8111(00)00061-8)
- Hemphill, L., & Roback, A. J. (2014). Tweet acts: How constituents lobby Congress via Twitter. In *Proceedings of the 17th ACM conference on computer supported cooperative work & social computing*, CSCW'14, pp. 1200–1210, New York, NY, USA. Association for Computing Machinery.
- Hickok, M. (2021). Lessons learned from AI ethics principles for future actions. *AI and Ethics*, 1(1), 41–47. <https://doi.org/10.1007/s43681-020-00008-1>
- Hogan, J., & Howlett, M. (2015). Reflections on our understanding of policy paradigms and policy change. In J. Hogan & M. Howlett (Eds.), *Policy paradigms in theory and practice: Discourses, ideas and anomalies in public policy dynamics* (pp. 3–18). Palgrave Macmillan UK.
- Hollanders, D., & Vliegthart, R. (2008). Telling what yesterday's news might be tomorrow: Modeling media dynamics. 33(1), 47–68. <https://doi.org/10.1515/COMMUN.2008.003>
- Hopkins, D. J., Kim, E., & Kim, S. (2017). Does newspaper coverage influence or reflect public perceptions of the economy? *Research & Politics*, 4(4), 1–7. <https://doi.org/10.1177/2053168017737900>
- Howlett, M. (1998). Predictable and unpredictable policy windows: Institutional and exogenous correlates of Canadian federal agenda-setting. *Canadian Journal of Political Science/Revue Canadienne de Science Politique*, 31(3), 495–524. <https://doi.org/10.1017/S0008423900009100>
- Imbrie, A., Dunham, J., Gelles, R., & Aiken, C. (2020). *Mainframes: A provisional analysis of rhetorical frames in AI*. Technical report, Center for Security and Emerging Technology.
- Imbrie, A., Gelles, R., Dunham, J., & Aiken, C. (2021). *Contending frames: Evaluating rhetorical dynamics in AI*. Technical report, Center for Security and Emerging Technology.
- Jones, B. D., & Baumgartner, F. R. (2004). Representation and agenda setting. *Policy Studies Journal*, 32(1), 1–24. <https://doi.org/10.1111/j.0190-292X.2004.00050.x>
- Jones, B. D., Larsen-Price, H., & Wilkerson, J. (2009). Representation and American governing institutions. *The Journal of Politics*, 71(1), 277–290. <https://doi.org/10.1017/S002238160809018X>
- Kingdon, J. W. (1984). *Agendas, alternatives, and public policies* (1st ed.). Little, Brown.
- König, P. D., & Wenzelburger, G. (2021). The legitimacy gap of algorithmic decision-making in the public sector: Why it arises and how to address it. *Technology in Society*, 67, 101688. <https://doi.org/10.1016/j.techsoc.2021.101688>
- Kuhlmann, S., & Rip, A. (2018). Next-generation innovation policy and grand challenges. *Science and Public Policy*, 45(4), 448–454. <https://doi.org/10.1093/scipol/scy011>
- Leung, J. (2020). *Who will govern artificial intelligence?* PhD thesis, Oxford University.
- Liu, X., Lindquist, E., & Vedlitz, A. (2011). Explaining media and Congressional attention to global climate change, 1969–2005: An empirical test of agenda-setting theory. *Political Research Quarterly*, 64(2), 405–419. <https://doi.org/10.1177/1065912909346744>
- Lodge, M., & Wegrich, K. (2015). Crowdsourcing and regulatory reviews: A new way of challenging red tape in British government? *Regulation & Governance*, 9(1), 30–46. <https://doi.org/10.1111/rego.12048>
- Lowi, T. J. (1964). American business, public policy, case-studies, and political theory. *World Politics*, 16(4), 677–715. <https://doi.org/10.2307/2009452>
- Macnaghten, P., & Chilvers, J. (2014). The future of science governance: Publics, policies, practices. *Environment and Planning C: Government and Policy*, 32(3), 530–548. <https://doi.org/10.1068/c1245j>
- McCombs, M. E. (2004). *Setting the agenda: The mass media and public opinion*. Polity; Blackwell Pub.
- McCombs, M. E., & Shaw, D. L. (1972). The agenda-setting function of mass media. *Public Opinion Quarterly*, 36(2), 176–187. <https://doi.org/10.1086/267990>

- McKenzie, E., & Gardner, E. S. (2010). Damped trend exponential smoothing: A modelling viewpoint. *International Journal of Forecasting*, 26(4), 661–665. <https://doi.org/10.1016/j.ijforecast.2009.07.001>
- Mellon, J., & Prosser, C. (2017). Twitter and Facebook are not representative of the general population: Political attitudes and demographics of British social media users. *Research & Politics*, 4(3), 2053168017720008. <https://doi.org/10.1177/2053168017720008>
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems*, 26.
- Minkinen, M., Zimmer, M. P., & Mäntymäki, M. (2023). Co-shaping an ecosystem for responsible AI: Five types of expectation work in response to a technological frame. *Information Systems Frontiers*, 25, 103–121. <https://doi.org/10.1007/s10796-022-10269-2>
- Mislove, A., Lehmann, S., Ahn, -Y.-Y., Onnela, J.-P., & Rosenquist, J. (2011). Understanding the demographics of Twitter users. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1), 554–557. <https://doi.org/10.1609/icwsm.v5i1.14168>
- Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2019). From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices. *Science and Engineering Ethics*, 26(4), 2141–2168. <https://doi.org/10.1007/s11948-019-00165-5>
- Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2023). Operationalising AI ethics: Barriers, enablers and next steps. *AI and Society*, 38, 411–423. <https://doi.org/10.1007/s00146-021-01308-8>
- Neuman, W. R., Neuman, R. W., Just, M. R., & Crigler, A. N. (1992). *Common knowledge: News and the construction of political meaning*. University of Chicago Press.
- O'Shaughnessy, M. R., Schiff, D. S., Varshney, L. R., Rozell, C. J., & Davenport, M. A. (2023). What governs attitudes toward artificial intelligence adoption and governance? *Science and Public Policy*, 50(2), 161–176. <https://doi.org/10.1093/scipol/scac056>
- Ouchchy, L., Coin, A., & Dubljević, V. (2020). AI in the headlines: The portrayal of the ethical issues of artificial intelligence in the media. *AI and Society*, 35(4), 927–936. <https://doi.org/10.1007/s00146-020-00965-5>
- Peeters, J., Van Aelst, P., & Praet, S. (2021). Party ownership or individual specialization? A comparison of politicians' individual issue attention across three different agendas. *Party Politics*, 27(4), 692–703. <https://doi.org/10.1177/1354068819881639>
- Peters, B. G., & Hogwood, B. W. (1985). In search of the issue-attention cycle. *The Journal of Politics*, 47(1), 238–253. <https://doi.org/10.2307/2131074>
- Pierre, J. (ed). (2000). *Debating governance: Authority, steering, and democracy*. OUP Oxford.
- Prates, M., Avelar, P., & Lamb, L. C. (2018). On quantifying and understanding the role of ethics in AI research: A historical account of flagship conferences and journals. *arXiv:1809.08328[cs]*, 188–173.
- Rein, M., & Schön, D. (1996). Frame-critical policy analysis and frame-reflective policy practice. *Knowledge and Policy*, 9(1), 85–104. <https://doi.org/10.1007/BF02832223>
- Ringel, S. (2021). Digitizing the paper of record: Archiving digital newspapers at the New York Times. *Journalism*, 24, 245–261. <https://doi.org/10.1177/14648849211023849>
- Ripberger, J. T. (2011). Capturing curiosity: Using internet search trends to measure public attentiveness. *Policy Studies Journal*, 39(2), 239–259. <https://doi.org/10.1111/j.1541-0072.2011.00406.x>
- Rowe, G., & Frewer, L. J. (2000). Public participation methods: A framework for evaluation. *Science, Technology, & Human Values*, 25(1), 3–29. <https://doi.org/10.1177/016224390002500101>
- Schiff, D., Borenstein, J., Laas, K., & Biddle, J. (2021). AI ethics in the public, private, and NGO sectors: A review of a global document collection. *IEEE Transactions on Technology and Society*, 2(1), 31–42.
- Shapiro, M. A., & Hemphill, L. (2017). Politicians and the policy agenda: Does use of Twitter by the U.S. Congress direct New York Times content? *Policy & Internet*, 9(1), 109–132. <https://doi.org/10.1002/po13.120>
- Shapiro, M. A., Hemphill, L., & Otterbacher, J. (2018). Updates to congressional speech acts on Twitter. In *American Political Science Association annual conference* Boston, MA, p. 23,.
- Sharp, E. B. (1994). The dynamics of issue expansion: Cases from disability rights and fetal research controversy. *The Journal of Politics*, 56(4), 919–939. <https://doi.org/10.2307/2132067>
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press.
- Shoemaker, P. J., & Reese, S. D. (1991). *Mediating the message: Theories of influences on mass media content*. Longman.
- Siddique, N. (2019). Creating a data set for 116th Congress Twitter handles. <https://ws-dl.blogspot.com/2019/04/2019-04-01-creating-data-set-for-116th.html>.

- Sloane, M., Moss, E., Awomolo, O., & Forlano, L. (2022). Participation Is not a Design Fix for Machine Learning. In Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, pp. 1–6. <https://doi.org/10.1145/3551624.3555285>
- Smits, R., & Kuhlmann, S. (2004). The rise of systemic instruments in innovation policy. *International Journal of Foresight and Innovation Policy*, 1(1/2), 4. <https://doi.org/10.1504/IJFIP.2004.004621>
- Soete, L. (2007). From industrial to innovation policy. *Journal of Industry, Competition and Trade*, 7(3-4), 273–284. <https://doi.org/10.1007/s10842-007-0019-5>
- Stahl, B. C. (2021). Addressing ethical issues in AI. In B. C. Stahl, Ed., *Artificial intelligence for a better future: An ecosystem perspective on the ethics of AI and emerging digital technologies*. (pp. 55–79). Springer International Publishing; SpringerBriefs in Research and Innovation Governance.
- Stark, L., Greene, D., & Hoffmann, A. L. (2021). Critical perspectives on governance mechanisms for AI/ML systems. In J. Roberge & M. Castelle (Eds.), *The cultural life of machine learning: An incursion into critical AI studies* (pp. 257–280). Springer International Publishing.
- Stirling, A. (2008). “Opening up” and “closing down”: Power, participation, and pluralism in the social appraisal of technology. *Science, Technology, & Human Values*, 33(2), 262–294. <https://doi.org/10.1177/0162243907311265>
- Stix, C., & Maas, M. (2021). Bridging the gap: The case for an ‘incompletely theorized agreement’ on AI policy. *AI and Ethics*, 1(3), 261–271. <https://doi.org/10.1007/s43681-020-00037-w>
- Stoker, K. (2014). Defining public in public relations: How the 1920s debate over public opinion influenced early philosophies of public relations. In: St. John III B., Lamme M. O. and L’Etang J. (Eds.), *Pathways to public relations*. Routledge.
- Taeihagh, A. (2021). Governance of artificial intelligence. *Policy and Society*, 40(2), 137–157. <https://doi.org/10.1080/14494035.2021.1928377>
- Tromble, R. (2018). The great leveler? Comparing citizen–politician Twitter engagement across three western democracies. *European Political Science*, 17(2), 223–239. <https://doi.org/10.1057/s41304-016-0022-6>
- Ulnicane, I. (2016). Grand challenges’ concept: A return of the ‘big ideas’ in science, technology and innovation policy? *International Journal of Foresight and Innovation Policy*, 11(1-3), 5–21. <https://doi.org/10.1504/IJFIP.2016.078378>
- Ulnicane, I. (2023). Against the new space race: Global AI competition and cooperation for people. *AI & Society*, 38, 681–683. <https://doi.org/10.1007/s00146-022-01423-0>
- Ulnicane, I., Eke, D. O., Knight, W., Ogoh, G., & Stahl, B. C. (2021). Good governance as a response to discontents? Déjà vu, or lessons for AI from other emerging technologies. *Interdisciplinary Science Reviews*, 46(1-2), 71–93. <https://doi.org/10.1080/03080188.2020.1840220>
- Ulnicane, I., Knight, W., Leach, T., Stahl, B. C., & Wanjiku, W.-G. (2020). Framing governance for a contested emerging technology: Insights from AI policy. *Policy and Society*, 40, 158–177. <https://doi.org/10.1080/14494035.2020.1855800>
- Vaccari, C., Valeriani, A., Barberá, P., Bonneau, R., Jost, J. T., Nagler, J., & Tucker, J. A. (2015). Political expression and action on social media: Exploring the relationship between lower- and higher-threshold political activities among Twitter users in Italy. *Journal of Computer-Mediated Communication*, 20(2), 221–239. <https://doi.org/10.1111/jcc4.12108>
- Vaida, F., & Blanchard, S. (2005). Conditional Akaike Information for mixed-effects models. *Biometrika*, 92(2), 351–370. <https://doi.org/10.1093/biomet/92.2.351>
- Van Aelst, P., & Walgrave, S. (2016). Information and arena: The dual function of the news media for political elites. *Journal of Communication*, 66(3), 496–518. <https://doi.org/10.1111/jcom.12229>
- Vesnic-Alujevic, L., Nascimento, S., & Pólvara, A. (2020). Societal and ethical impacts of artificial intelligence: Critical notes on European policy frameworks. *Telecommunications Policy*, 44(6), 101961. <https://doi.org/10.1016/j.telpol.2020.101961>
- Vliegenthart, R. (2014). Moving up. Applying aggregate level time series analysis in the study of media coverage. *Quality and Quantity*, 48(5), 2427–2445. <https://doi.org/10.1007/s11135-013-9899-0>
- Vliegenthart, R., & Walgrave, S. (2008). The contingency of intermedia agenda setting: A longitudinal study in Belgium. *Journalism & Mass Communication Quarterly*, 85(4), 860–877. <https://doi.org/10.1177/107769900808500409>
- von Schomberg, R. (2013). A vision of responsible research and innovation. In: Owen R., Bessant J. and Heintz M. (Eds.), *Responsible innovation* (pp. 51–74). Wiley-Blackwell.

- Warnke, P., Koschätzky, K., Dönitz, E., Zenker, A., Stahlecker, T., Som, O., Cuhls, K., & Güth, S. (2016). *Opening up the innovation system framework towards new actors and institutions*. Working Paper 49, Fraunhofer ISI Discussion Papers - Innovation Systems and Policy Analysis.
- Weaver, D. A., & Bimber, B. (2008). Finding news stories: A comparison of searches using LexisNexis and Google News. *Journalism & Mass Communication Quarterly*, 85(3), 515–530. <https://doi.org/10.1177/107769900808500303>
- White House. (2020). *Promoting the use of trustworthy artificial intelligence in the federal government*. (Executive Order No. 13960). <https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government>.
- Wilkerson, J., & Casas, A. (2017). Large-scale computerized text analysis in political science: Opportunities and challenges. *Annual Review of Political Science*, 20(1), 529–544. <https://doi.org/10.1146/annurev-polisci-052615-025542>
- Wolfe, M., Jones, B. D., & Baumgartner, F. R. (2013). A failure to communicate: agenda setting in media and policy studies. *Political Communication*, 30(2), 175–192. <https://doi.org/10.1080/10584609.2012.737419>
- Wood, B. D., & Peake, J. S. (1998). The dynamics of foreign policy agenda setting. *American Political Science Review*, 92(1), 173–184. <https://doi.org/10.2307/2585936>
- Yanovitzky, I. (2002). Effects of news coverage on policy attention and actions: A closer look into the media-policy connection. *Communication Research*, 29(4), 422–451. <https://doi.org/10.1177/0093650202029004003>
- Zaller, J. (1992). *The nature and origins of mass opinion*. Cambridge University Press.
- Zhang, D., Saurabh, M., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., Sellitto, M., Shoham, Y., Clark, J., & Perrault, R. (2021). *Artificial Intelligence Index report 2021*. Technical report, AI Index Steering Committee, Human-Centered AI Institute, Stanford University.