

Looking through a policy window with tinted glasses: Setting the agenda for U.S. AI policy

Daniel S. Schiff 

Department of Political Science, Purdue University, West Lafayette, Indiana, USA

Correspondence

Daniel S. Schiff, Department of Political Science, Purdue University, West Lafayette, IN, USA.
Email: dschiff@purdue.edu

Abstract

The policy agenda is currently being established for artificial intelligence (AI), a domain marked by complex and sweeping implications for economic transformation tempered by concerns about social and ethical risks. This article reviews the United States national AI policy strategy through extensive qualitative and quantitative content analysis of 63 strategic AI policy documents curated by the federal government between 2016 and 2020. Drawing on a prominent theory of agenda setting, the Multiple Streams Framework, and in light of competing paradigms of technology policy, this article reviews how the U.S. government understands the key policy problems, solutions, and issue frames associated with AI. Findings indicate minimal attention to focusing events or problem indicators emphasizing social and ethical concerns, as opposed to economic and geopolitical ones. Further, broad statements noting ethical dimensions of AI often fail to translate into specific policy solutions, which may be explained by a lack of technical feasibility or value acceptability of ethics-related policy solutions, along with institutional constraints for agencies in specific policy sectors. Finally, despite widespread calls for increased public participation, proposed solutions remain expert dominated. Overall, while the emerging U.S. AI policy agenda reflects a striking level of attention to ethics—a promising development for policy stakeholders invested in AI

A version of this article is also part of a doctoral dissertation.

ethics and more socially oriented approaches to technology governance—this success is only partial and is ultimately layered into a traditional strategic approach to innovation policy.

KEYWORDS

agenda-setting, artificial intelligence, national strategy, social and ethical implications of technology, technology governance

INTRODUCTION

Artificial intelligence (AI) policy is at a nascent stage. While dozens of countries have begun to put forward preliminary policy strategies, relatively few have adopted formal regulation (CIFAR, 2020; OECD, 2021). Indeed, the 2020s are expected to be a period of significant policymaking activity (Perry & Uuk, 2019; Zhang et al., 2022). As such, AI policy is arguably best understood to be undergoing agenda-setting, the often contentious process wherein policy actors compete to shape the key policy problems and solutions that will set the stage for future policymaking. Yet while policy process theory has been fruitfully applied to understand many policy domains, prominent agenda-setting frameworks like the Multiple Streams Approach or Framework (MSF) have rarely been applied to technology policy (Jones et al., 2016) and little is known about agenda-setting for AI policy in particular (Taeihagh, 2021), despite its sweeping social and economic implications.¹

As part of a Special Issue on AI Politics and Policy, this article seeks to extend scholarly knowledge on the drivers and tensions underpinning agenda-setting in AI policy, drawing on the MSF along with competing paradigms in innovation policy with disparate implications for technology governance. In particular, while traditional innovation policy often emphasizes strategic economic and geopolitical dimensions of technology underpinned by a strong expert orientation, newer paradigms emphasize societal objectives of innovation and call for broader public participation. How the AI policy agenda unfolds in this context helps to reveal the prospects for incorporating AI ethics into policymaking, which actors have power in the agenda-setting process, and how global AI governance is likely to take form.

To provide insight on these issues, this study performs qualitative and quantitative content analysis of 63 key strategic AI documents identified by the United States (U.S.) federal government and published between 2016 and 2020. Through extensive coding and analysis, it examines the policy problems, solutions, broader issue frames, and key focusing events and indicators reflected in U.S. AI policy discourse, as well as the role of experts and the general public in this process. Drawing on these data thus helps to unpack aspects of the policy process theorized as important in shaping the AI policy agenda. Based on this analysis, this study engages with the following question:

- Does the U.S. federal AI policy agenda better reflect a traditional approach to innovation policy, emphasizing economic and geopolitical goals, or a ‘transformative’ paradigm, emphasizing social and ethical objectives and a public rather than expert orientation?

Results indicate, on one hand, that advocates of AI ethics² have had remarkable success in promoting attention to social and ethical dimensions of AI in policy discourse generally. Indeed,

the degree of attention to ethics in technology governance may be unprecedented. Yet, such a transformation is partial at best and substantially tempered in practice. For instance, despite broad-based calls for attention to ethical implications at the mission statement level, the overwhelming majority of focusing events and indicators discussed in U.S. AI policy documents are traditional in nature, emphasizing economic benefits of AI and associated geopolitical concerns. Further, the majority of policy problems and solutions featured in these documents are also traditional in nature, including calls for increased research and development to realize AI's benefits, responses to military and security implications, expansion of access to datasets to support AI development, and growth of the highly skilled technology workforce through increased education. While some social and ethical problems do receive significant attention, such as those surrounding privacy, transparency, and trust, even these often take on a 'hybrid' mode where associated policy solutions are interpreted and justified in light of realizing traditional innovation goals, such as promoting trustworthy AI to increase consumer adoption. Meanwhile, other ethical concerns such as fairness, inequality, and human rights receive somewhat less attention. Finally, despite almost ubiquitous calls for public and diverse participation, the large majority of concrete recommendations specify industry or government experts, undermining the prospects for true public governance of AI.

The study also considers potential and preliminary explanations for these findings, drawing on the MSF and innovation policy literature. First, some ethical problems are thought to be addressable through technical fixes, and thus more concrete proposals surrounding topics like privacy may benefit from the ostensible 'technical feasibility' of these solutions. Meanwhile, broader calls for societal transformation that might challenge innovation-oriented goals, such as addressing inequality, may run afoul of 'value acceptability' in the U.S. context. Third, high-level rhetorical interest in social and ethical implications of AI, especially in documents with a government-wide scope, may fail to translate into action as government agencies in specific sectors interpret these questions through traditional norms and policy instruments. Importantly, these explanations are not complete and need further work using a plurality of methods to explore further.

Given the importance of the United States in global technology governance as well as for AI research and development, these findings are important to international stakeholders interested in AI policy, as well as scholars of policy process theory and innovation policy interested in the extent to which new technology policy domains are characterized by more participatory, societal objectives and processes. Ultimately, despite some striking success for stakeholders invested in AI's social and ethical implications, the translation of ethics into policy appears currently limited. Instead of reflecting a wholesale evolution from traditional innovation policy to more transformative notions of technology governance, the U.S. AI policy agenda instead reflects a layering or subsumption of the latter into the former. Nevertheless, policy entrepreneurs who are better able to translate their social and ethical concerns into concrete policy solutions may find a willing audience. While the policy agenda is taking shape, the policy window is not yet closed.

The paper begins by presenting its conceptual framework on agenda-setting, the Multiple Streams Framework, and the contest between traditional and transformative notions of innovation policy. It then reviews the paper's dataset of federal AI policy documents along with the codebook and mixed methods analysis approach. The results describe patterns and potential explanations related to the conceptual framework, covering focusing events and indicators, issue frames, policy problems and solutions, and the role of experts and the public. Finally, the discussion and conclusion reflect on possible explanations, limitations, and what the study's findings might imply for the governance of AI and role of ethics in technology policy going forward.

THEORETICAL BACKGROUND

Agenda-setting and the multiple streams framework

The MSF is the first of two theoretical lenses used to study the policy context and agenda-setting process surrounding AI. Originally drawing on the concept of “organized anarchies,” (Cohen et al., 1972), but applied to institutional rather than organizational settings, the MSF’s basic assumptions involve ambiguity, unclear technology, problematic preferences, time constraints, and fluid participation (Cairney & Jones, 2016; Herweg et al., 2017). In short, the agenda-setting process is understood to be complex, characterized by ambiguity in issues, processes, and participation, leading to similar complexity in formulating policy preferences under time and attentional constraints. As such, competition over ideas plays a key role in how policy agendas are formulated (Greer, 2016), and constitutes a focus of this article.

The MSF is centered around three partially independent ‘streams’ of policy-relevant activity. The *problem stream* considers how conditions in society come to be considered as problems that demand policy action. Policy entrepreneurs and problem brokers (Knaggård, 2015) draw on various focusing events, quantitative indicators, and feedback from existing policy programs to identify and define (or indeed, *construct*) the relevant problems. Meanwhile, relatively loose networks of often less visible actors work to develop policy alternatives or solutions in the *policy (or solution) stream*. Preliminary ideas, part of a “primeval soup” (Kingdon, 1984), are then refined through a “softening” process by policy communities with relevant expertise (Herweg, 2016) according to criteria such as technical and financial feasibility and value acceptability. Finally, in the *political stream*, national mood or public opinion, interest group activity, and the composition of the current government (Herweg et al., 2015) shape receptivity to the proposed policy problems and solutions. Here, policy entrepreneurs play a key role in coupling the streams together to place a package on the policy agenda (Mintrom & Norman, 2009; Roberts & King, 1991). This article highlights a subset of theoretical elements deemed important by the MSF, and uses these to guide the empirical approach in light of the policy context and data analyzed.³

Figure 1 provides an overview of elements from the MSF and their role in agenda-setting for AI policy, augmented by the subsequent discussion on competing policy paradigms. Key elements theorized or determined to be important are emphasized.

Competing paradigms of technology policy: Traditional and transformative

The second conceptual pillar for the case study is the contrast between competing paradigms of innovation (or technology) policy. Here, a paradigm is understood to encompass the intersubjectively held normative goals behind policymaking, conceptions of underlying policy problems and appropriate instruments used to achieve these goals, and the underlying terminology and ideas used to frame this process (Daigneault, 2014; Hall, 1993). Indeed, policy paradigms have a storied history in the realm of technology and innovation policy (Morlacchi & Martin, 2009). Particularly relevant to this study is the extent to which policy paradigm shifts involve radical breaks in the Kuhnian sense, or more evolutionary transitions within the bounds of ‘normal’ policymaking (Princen & ‘t Hart, 2014) as part of the sociological and ideational process of change. To assess this process in the context of AI policy, this study seeks to provide “direct evidence of policy actors’ ideas and beliefs” (Daigneault, 2014, p. 463) including measurable evidence of the “power

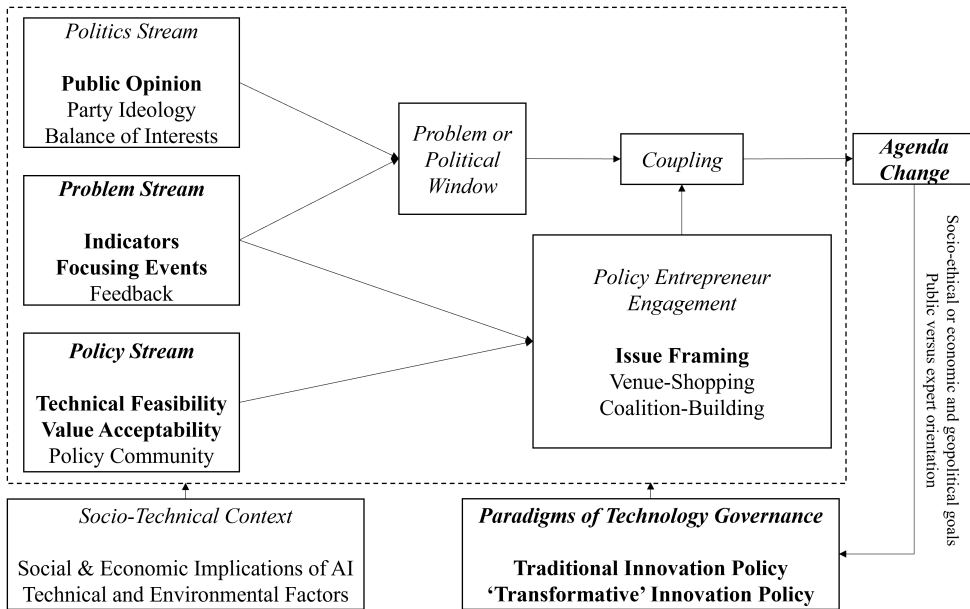


FIGURE 1 Conceptual framework: Elements of agenda-setting in AI policy.

and salience of the underlying ideas” (Baumgartner, 2014, p. 478) through quantifying attention to policy problems, solutions, and underlying issue frames.

In particular, the first paradigm this study considers is termed the ‘traditional’ approach to innovation policy, in which an economic, firm-centered, and expert-driven logic dominates agenda-setting and policy design (Edler & Fagerberg, 2017). For such an approach, strategies aimed at boosting the workforce, fostering economic growth and innovation, and accelerating adoption of emerging technologies are paramount. Departing somewhat from traditional industrial policy, modern innovation policy is especially focused on high-tech industries and technological advances that may lead to growth acceleration (Soete, 2007). Associated with this logic is an issue frame which emphasizes the economic growth potential associated with AI in terms of increased productivity, efficiency, industrial profit, national GDP, entrepreneurial activity, and so on (Imbrie et al., 2021; Ulicane et al., 2022).

Typical strategies aligned with traditional innovation policy thus often center on science, technology, engineering, and math (STEM) and associated supply-side reforms, particularly expanding the size and skill of the workforce through education, federal research and development funding, and export and investment controls (Atkinson & Mayo, 2010; Fischer et al., 2021). Other innovation instruments include empowering regions, clusters, or hubs of innovation, supporting small- and medium-sized enterprises (SMEs), and aligning academic, industry, and public sector activity (Smits & Kuhlmann, 2004). Reference to the innovation frame is pervasive throughout policy discourse including in national AI policy strategies and key U.S. policy documents (Schiff, 2022).

A second and competing paradigm of technology policy deviates from a more economic, expert-driven, and firm-centered logic, emphasizing instead societal objectives and public participation in policy. Such an approach is in part a response to failures of traditional science and technology policy and innovation systems policy to solve major societal challenges such as inequality and climate change (Ulicane, 2016). This perspective can be understood as

associated with movements such as transformative innovation policy (Grillitsch et al., 2018; Schot & Steinmueller, 2018), mission-oriented innovation policy (Mazzucato, 2018) and Responsible Innovation (de Saille, 2015), jointly representing a “more recent shift in technology policy towards societal objectives” (Ulnicane et al., 2021, p. 80). Proponents of this paradigm are arguably more prone to advocate for stricter regulation and precaution in policymaking in recognition of the social and ethical risks of AI (Owen et al., 2013; Stirling, 2016). This approach thus acknowledges the possibility that innovation can lead to ethical and societal harms, such that a ubiquitous pro-innovation bias may be problematic. Relatedly, it tends to avoid a narrow focus on supply-side strategies aimed at economic growth, and embraces broader societal objectives across multiple sectors, potentially with a mission-oriented and global scale (Kuhlmann & Rip, 2018). For simplicity, this study terms this the ‘transformative’ paradigm, but does not treat it as synonymous with transformative innovation policy (Diercks et al., 2019).

This approach to technology governance also implies and typically explicitly recommends a more inclusive and participatory process in shaping the policy agenda, for which it is important to incorporate a broader array of actors beyond firms, academia, and government, such as civil society and the public (Warnke et al., 2016). Already, there have been repeated and widespread calls for diverse and public participation in AI policy from governments, academics, civil society, and industry (Schiff, Borenstein, et al., 2021; Stix, 2021; Ulnicane et al., 2020; Vesnic-Alujevic et al., 2020).⁴ Indeed, there is increased attention to the role of the public in policy discourse generally (Rowe & Frewer, 2000), within science and technology policy specifically (Macnaghten & Chilvers, 2014; Stirling, 2008), and even within AI policy (Buhmann & Fieseler, 2022; Stark et al., 2021). Yet, while this focus on public participation is emphasized in current technology and AI policy discourse, these calls also reflect the renewal of arguments in the 21st century for multi-stakeholder governance of emerging technologies in contrast to traditional top-down government (Pierre, 2000), and indeed follow a long history of political thought on the value of public participation (Arnstein, 1969; Dahl, 1978). Evidently, advocates for more public involvement do not feel their ambition has been achieved, a concern this study reiterates.

In summary, an important scholarly and policy question is therefore whether technology governance is indeed characterized by such a shift. The MSF is valuable in this effort toward thinking more deeply about the underlying policy problems, appropriate solutions, roles of stakeholders, and goals behind policymaking. Along these lines, Table 1 draws on the specified components of the MSF and concept of competing paradigms, and presents competing, provisional predictions (or rough hypotheses). These derived predictions establish preliminary analytical grounds upon which to evaluate the research question.⁵

METHODOLOGY

Overview and epistemology

The methodology for this chapter, driving both data selection and analysis, is the single, holistic case study approach (Yin, 2018) realized through qualitative and quantitative document and content analysis. The unit of analysis is the *agenda-setting process* as reflected in 63 key AI policy documents curated by the U.S. federal government and published between January 2016 and December 2020. The rationale for the selection of this case is that AI policy is a new, revelatory, and empirically important emerging technology policy domain. Given clear evidence that the AI agenda-setting discourse reflects an unusual degree of contestation between various economic,

TABLE 1 Competing predictions for AI policy based on elements of the MSF

	Traditional paradigm	Transformative paradigm
Key indicators and focusing events	Key indicators and focusing events discussed in the U.S. AI policy agenda surround strategic economic and international competition objectives and concerns. They are less salient, accessible, and related to public concerns, and reflect an expert orientation	Key indicators and focusing events discussed in the U.S. AI policy agenda surround societal and ethical objectives and concerns. They are relatively more salient, accessible, and related to public concerns
Policy problems, solutions, and issue frames	Policy problems, solutions, and frames discussed in the U.S. AI policy agenda emphasize economic objectives, such as supply-side strategies for STEM education and R&D, targeted at economic growth and national competitiveness. A pro-innovation bias is present	Policy problems, solutions, and frames discussed in the U.S. AI policy agenda emphasize societal and ethical objectives, such as reducing inequality and serving vulnerable groups. Innovation is viewed more skeptically and with more precaution advocated
Role of the public and experts	Experts, acting as policy entrepreneurs or policy community members, play an outsized role in shaping policy problems, solutions, and frames discussed in the U.S. AI policy agenda. The need for technical expertise, and increasing the access, supply, and training of experts are emphasized	Members of the broader public play an outsized role in shaping policy problems, solutions, and frames discussed in the U.S. AI policy agenda. The need for diverse and inclusive participation, increased access, and public engagement and opinion in policy are emphasized

social, and ethical goals, because robust public participation has been called for extensively, and because this case was previously inaccessible to study, it is arguably a strong critical test of whether technology policy (and governance more broadly) is evolving. In short, if such a shift is likely to take place in technology policy, it should be in this domain if nowhere else.

A case study of the U.S. federal AI policy agenda in particular is valuable for several reasons. First, it promotes structured exploration of AI agenda-setting in a scholarly and policy context otherwise marked by diverse and sometimes piecemeal interdisciplinary conversation (Taeihagh, 2021). As such, it offers the possibility of integrating disparate knowledge on the conditions, influences, and actors relevant to AI agenda-setting using a prominent framework and detailed empirical evidence. Further, it offers descriptive, explanatory, and exploratory insights and establishes a helpful foundation for identifying more specific and future research questions with respect to the substantive and theoretical context studied in the article.

The ontological and epistemological approach employed here is most akin to critical realism. Critical realism acknowledges the existence of objective entities and mechanisms working in the world, but appreciates that these are filtered through human experience and subjectivity (Bhaskar, 1979; Wynn & Williams, 2012). In short, the observable world is understood as theory-laden, but not theory-determined (Fletcher, 2017; Hoddy, 2019). This is appropriate for the context of agenda-setting, given how objective features of the socioeconomic context and of

AI as a set of technologies are filtered through policy images, narratives, issue frames, and so on as part of the deliberative process. This broader reality can only be indirectly accessed, in this case by interpreting the discourse present in the U.S. AI policy documents.

Data and inclusion criteria

The dataset is composed of 63 “key strategic documents” for AI that apply at the federal or national level and are hosted at [AI.gov](https://ai.gov) (National AI Initiative Office, Office of Science and Technology Policy, 2022).⁶ These include a diversity of document types, including budgetary documents, requests for information, scientific and technical reports, strategy documents, ethical principles, and international agreements. Correspondingly, these documents are produced by a variety of authors, ranging from actors within the U.S. White House and Congressional Research Service to agencies such as the Departments of State and Defense. Some documents are initiated based on calls from actors in other *areas* of government, while others emerge from concerns within the specific organizations themselves. For example, many documents are explicit responses to Executive Orders and the American AI Initiative.

Yet, while documents have somewhat different goals, tones, and ultimately roles within the agenda-setting process, the large majority of documents are motivated in a similar fashion, either implicitly or explicitly. The overall motivation that connects the documents is a goal to better understand the implications of AI related to one’s purview, and to consider responses to these implications. Thus documents typically included both analytical and prescriptive elements, and this similarity in function was especially true for the most common document types: reports, requests for information, and strategy documents.⁷

Importantly, as these documents were explicitly curated by the federal government to reflect AI policy priorities, they serve as a logical starting place to understand the deliberative process influencing the emerging AI policy agenda. These official policy documents can thus be understood as “vehicles of messages, communicating or reflecting official intentions, objectives, commitments, proposals, ‘thinking’, ideology and responses to external events” (Freeman & Maybin, 2011, p. 157). Moreover, the content therein constitute “the outcome of a political process” influenced by multiple actors with competing discourses, and thus reflect the developing perspectives and preferences of policy actors over different alternatives with which they have been presented (Diercks et al., 2019, p. 887). As such, these documents are especially helpful for exploring to what extent certain problems, solutions, couplings, indicators of national mood, and so on, are influential in shaping the policy agenda.

Given the intentional choice for the government to curate these specific documents, this study approaches inclusion/exclusion with strong deference to inclusion in the final analytical dataset. However, the final sample excludes three documents from the original 66 over the specified time period of 2016–2020: one document is a duplicate and two contain very little discussion of AI.⁸ Figure 2 presents basic information about the dataset while Table S1.1 in the Appendix includes the complete list of analyzed documents.

An important feature of this analysis is that the dataset is centered on executive branch documents, and thus reflects only a subset of the broader agenda-setting process. For example, legislative hearings, private one-on-one meetings, civil society or private sector lobbying efforts, and legislator communications to constituents are not studied apart from possible mentions in the executive agency documents. As such, only some of the key actors, institutions, and processes that constitute the totality of agenda-setting are examined here, and many of the visible and




FIGURE 2 Analytical sample: U.S. federal AI policy documents (2016–2020). Author's calculations from AI.gov, $n = 63$.

invisible dynamics require additional research to understand. This scope condition implies limitations with respect to the completeness and generalizability of the findings.

However, there are several key reasons for viewing these documents as a valuable source to understand agenda-setting. First, while document text is often ‘strategic’ in the sense that underlying attitudes are at least partially hidden, the executive branch of government is thought to be more neutral in its approach than the legislature (Miller, 2015; Stivers, 2015). Second, the executive branch plays a critical role not only in informing the preliminary stages of agenda-setting by producing analyses often explicitly requested by other policymakers; but also in ultimately shaping and implementing policy. Finally, policy problems and solutions described via lengthy technical reports, for example, are more likely to have the necessary amount of detail to make inferences with confidence as compared to legislative stump speeches. Thus, while it is difficult to understand the complete set of actors and documents involved in AI policy agenda-setting in the context of a single study, these (primarily) executive agency documents are amongst the best sources to understand the issues studied here.

Analysis approach

The analysis relies initially on an a priori (or directed) coding methodology, drawing from the literature on AI policy and the article’s conceptual framework to create an initial codebook (Hsieh & Shannon, 2005; Miles et al., 2013). For instance, the codebook includes a set of high-level domains and sub-codes articulating specific AI policy problems, policy solutions, and issue frames, as well as focusing events, indicators of problems, stakeholders, and other areas.⁹ Importantly however, the initial codebook was substantially modified, with additions, deletions, along with merging and separation of existing codes, and specification of more robust definitions (Hoddy, 2019).¹⁰

TABLE 2 Final codebook

Coding domains (# of codes)	Code
Focusing Events (7)	Traditional: <ul style="list-style-type: none"> • Expert Concerns • Games • Geopolitics & Military • Industry & Government Adoption • Technical Performance & Advances Transformative: <ul style="list-style-type: none"> • Protests • Scandals & Disasters
Indicators (7)	Traditional: <ul style="list-style-type: none"> • Economic & Workforce • Education • Expert Concerns • Geopolitics & Military • Technical Performance & Advances Transformative: <ul style="list-style-type: none"> • Poverty, Harm, & Fairness • Scandals & Disasters
Issue Frames (3)	Traditional: <ul style="list-style-type: none"> • Geopolitics • Innovation Transformative: <ul style="list-style-type: none"> • Ethics
Problems (16)	Traditional: <ul style="list-style-type: none"> • Data Quality & Access • Misuse & Hostile Actors • Performance & Reliability • Realizing Benefits • Security & Military • Workforce & Education Hybrid: <ul style="list-style-type: none"> • Accountability & Responsibility • Privacy • Safety • Transparency • Trust Transformative: <ul style="list-style-type: none"> • Fairness & Bias • Human & Civil Rights • Inequality & Inclusion • Value Alignment • Vulnerable Populations
Solutions (16)	Traditional: <ul style="list-style-type: none"> • Cooperation & Dialogue • Data Quality & Access • Pilot Projects & Testbeds • R&D & Adoption • Workforce & Education

(Continues)




TABLE 2 (Continued)

Coding domains (# of codes)	Code
	Hybrid: <ul style="list-style-type: none"> • Build Trust • Diverse Participation • Grants & Procurement • Human-AI Teaming • Impact & Risk Assessment • Monitoring & Reporting • Standards/Best Practices • Transparency Transformative: <ul style="list-style-type: none"> • Public Engagement • Social/Ethical Consideration • Technical Fixes for Ethics
Stakeholders (2)	<ul style="list-style-type: none"> • Experts • Public

Note: Note that while Birkland and Warnement (1998, 2016) define focusing events as sudden, relatively uncommon, typically harmful, and restricted to a particular location, the definition used here deviates, adopting a broader working definition to capture ‘events’ as reflected in AI policy discourse, often less discrete and region-specific. Along similar lines, policy problems and solutions are defined somewhat flexibly to be responsive to the nature of the emerging policy discourse.

Table 2 provides a summary of the key domains analyzed and Table S1.2 in the Appendix provides definitions and representative quotes for each code. Table 2 organizes codes by domain, and for some domains includes categories for ‘traditional,’ ‘transformative,’ and ‘hybrid,’ a classification approach explained in more detail in the Results section.

The coding process thus sought to answer questions like:

- Are focusing events and indicators referenced, and if so, which seem to play an important role?
- What are the key policy problems, solutions, and issue frames proposed?
- What kinds of actors are considered important in shaping the AI agenda (e.g., the public, experts)?
- To what extent do socio-ethical considerations translate into concrete policy solutions?

The application of the codebook, using qualitative analysis software Atlas.ti (version 22), involved identifying all relevant codes that applied at the level of individual paragraphs, helpful for capturing surrounding context and analyzing the co-occurrence of codes (e.g., of experts and particular policy solutions). Of course, this approach involves trade-offs, as documents vary widely and structural differences mean that the coding strategy is both under and overinclusive at times.¹¹ Indeed, while the documents (and individual instances of codes) are diverse in terms of style, tone, purpose, authorship, and importance, this diversity is a feature of policy discourse. As such, the approach taken here is based on the understanding that quantification, aggregation, and comparison are necessarily imperfect, and that the results represent a broad, heuristic overview of a multi-faceted discourse. Another limitation is that the codebook development and coding was a single-coder process,¹² which means that, for example, inter-rater reliability measures are not available. However, for this kind of qualitative analysis, other measures of rigor may be more appropriate (Lincoln & Guba, 1986), such as trustworthiness, authenticity,

and confirmability, facilitated by transparency in methodology and access to the underlying data sources and codebook. The coding resulted in over 4100 individual paragraphs coded, with codes applied to around 2300 pages of text across the 63 documents.

Following the coding process, the article draws on quantitative and qualitative content analysis techniques (White & Marsh, 2006). For example, to understand the importance of policy problems quantitatively, I extract statistics capturing the absolute prevalence of each problem (number of coded paragraphs per document and overall) and binary presence of each problem per document (present or absent). Importantly, I also determine the ‘normalized’ importance of each problem, by restricting attention to only policy problems (or solutions, or focusing events, etc.), and by calculating percentage representation of each problem per document. This allows me to rank and quantify the salience of given policy problems, solutions, issue frames, and so on, while addressing concerns like the possibility that very large documents with many codes or strong emphases would skew results in a given direction.¹³ This strategy in part—but certainly not in whole—buttresses the presentation of aggregate results given diverse documents.

Qualitatively, the analysis also involved creating daily memos during the coding process, which reported on possible findings, patterns, exceptions, and associations with the theoretical framework (Birks et al., 2008). I further engaged in analysis of quotations captured in single codes, codes within domains (e.g., all focusing events), and performed cross-code analysis (e.g., examining the intersection of particular policy stakeholders and policy solutions), guided by the conceptual elements of interest and in light of emerging potential explanations.

Finally, I applied an abductive (or retroductive) explanatory approach to synthesize quantitative and qualitative results. This involves employing pattern matching in light of the conceptual framework and provisional predictions (Bouncken et al., 2021; Sinkovics, 2018), and associating what is empirically observed with plausible generative mechanisms (Avenier & Thomas, 2015; Gioia et al., 2013). By redescribing the empirical data in this way, the study thus seeks to confront and elaborate theory to provide a more robust explanatory account of AI policy agenda-setting.

RESULTS

Focusing events and indicators: Dominance of traditional concerns

The first results address the relative attention paid to different types of focusing events and indicators, understood as key elements of the problem stream that lead societal conditions to be recognized as pressing policy problems that demand action (DeLeo et al., 2021). Each focusing event or indicator is classified as ‘traditional’ or ‘transformative’ as per Table 2.¹⁴ Figure 3 shows coverage of focusing events and indicators, arranged by normalized percentage.¹⁵ Overall, the large majority of focusing events and indicators mentioned in U.S. AI policy documents are traditional in nature. For example, while only 10 documents discuss at least one transformative focusing event, 40 mention at least one traditional focusing event. For indicators, similarly, 13 documents mention a transformative indicator compared to 31 which mention a traditional indicator. Nearly 96% of focusing events and 84% of indicators mentioned are traditional in nature.¹⁶

The slant of these results is to some degree surprising. Despite widespread attention to AI ethics topics in media generally (Chuan et al., 2019; Perrault et al., 2019), including ostensibly high-profile scandals surrounding biased facial recognition, wrongful arrests, sexism in hiring algorithms, self-driving car crashes, and more, very few of these incidents appear in U.S. AI policy documents. Further, despite prominent media and AI stakeholder discourse surrounding expres-

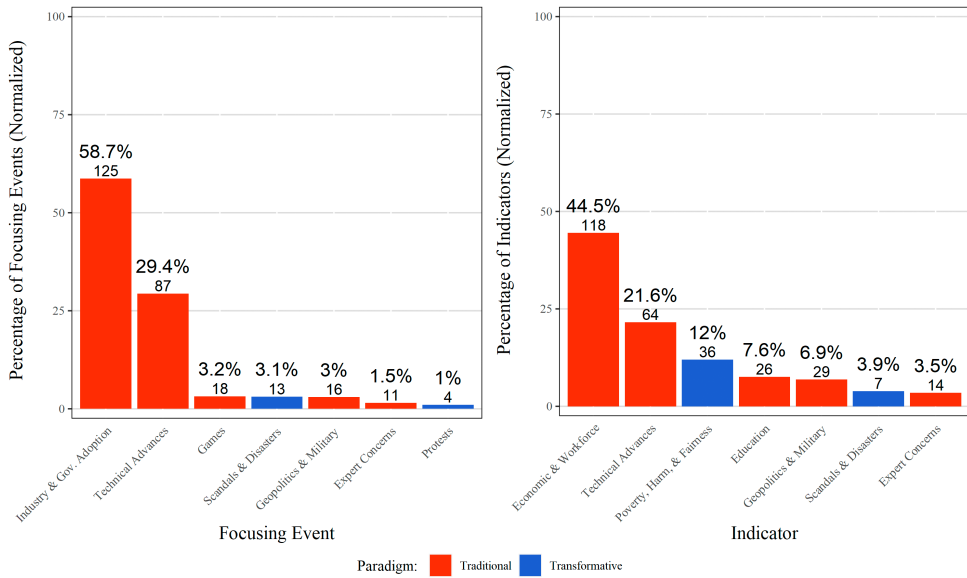


FIGURE 3 Coverage of focusing events and indicators in U.S. AI policy documents. Percentages of all focusing events (total $n = 274$) and indicators (total $n = 294$) in documents. Normalized percentages with absolute totals below.

sions of expert concern by individuals like Stephen Hawking and Elon Musk (Galanos, 2019; Neri & Cozman, 2020), these concerns are hardly present. Thus while there has been debate about the possibility that some ethics-related focusing events have been important (Ouchchy et al., 2020; Stix & Maas, 2021), this study suggests the impact of these influences is limited. Indeed, even recent efforts to catalog hundreds of incidents of social and ethical harm by the Partnership on AI (McGregor, 2020) in its new AI Incidents Database or to catalog protests via the Collective Action in Tech database (Tarnoff et al., 2021) are barely reflected in AI policy documents.

Instead, the dominance of traditional problem indicators comports with the focus of many global AI indicator dashboards that assess data infrastructure, research productivity, and human capital, such as the Global Cities AI Readiness Index, Global AI Index, and Global AI Vibrancy Index. While some indicator sets like the AI Social Contract Index, the OECD AI Policy Observatory database, and Government AI Readiness Index have more recently incorporated elements surround topics like ethics and diversity, the metrics overall are overwhelmingly traditional, an observation supported in Erkkilä’s recent extensive analysis of global metrics of AI (Erkkilä, 2023).

A possible explanation for this skew is that the social and ethical indicators and focusing events have influenced public opinion or the attitudes of legislators, but with less visible downstream implications, and that executive agencies are institutionally less prone to discuss these kinds of elements. Yet even in this charitable case, the finding that actors across more than a dozen sectors of government (many of which do articulate social and ethical concerns) overwhelmingly focus on indicators and focusing events surrounding the economy, workforce, and race for AI innovation generally is telling.

Issue frames: Synthesis with subsumption of ethics into innovation

This paper reviews discussion of three issue frames surrounding AI’s innovative, geopolitical, and ethical dimensions respectively, issues also found to be prominent in AI policy discourse by other researchers (Kim, 2023; Köstler & Ossewaarde, 2022; Ulnicane et al., 2022). Results indi-

cate that issue (or policy) frames surrounding both innovation and ethics feature heavily in U.S. AI policy documents. For instance, 59 of 63 documents discuss the innovation frame at some point, while almost as many documents, 57, discuss the ethics frame, with only 37 discussing the geopolitics frame. Moreover, the majority of documents discuss multiple frames. As displayed in Figure 4, 19 documents discuss both innovation and ethics, and a full 36 discuss all three issue frames. This provides compelling evidence that there has been some degree of convergence or synthesis in AI policy discourse, and indicates that these policy frames can be mutually reinforcing as opposed to mutually exclusive. In fact, some documents take very explicit efforts to justify the mutually reinforcing nature of these issues.

For example, conveying the overlap between geopolitics and ethics, the Congressional Research Service (CRS) (November 2020, p. 21) states that “it may be important for Congress to understand the state of rival AI development—particularly because U.S. competitors may have fewer moral, legal, or ethical qualms about developing military AI applications.” As an example of the merger of innovation and competition frames, a major piece of U.S. legislation is literally entitled the “U.S. Innovation & Competition Act” (2021). As stated by the National Science and Technology Council (NSTC) (September 2019, p. 3), “sustained Federal R&D investment in emerging technologies is critical to promoting and protecting American innovation and international leadership.” As an example of merging all three frames, the Organization for Economic Cooperation and Development (OECD) (May 2019, p. 6) argues that “trustworthiness of AI systems is a key factor for the diffusion and adoption of AI ... essential to fostering adoption of trustworthy AI in society, and to turning AI trustworthiness into a competitive parameter in the global marketplace.”

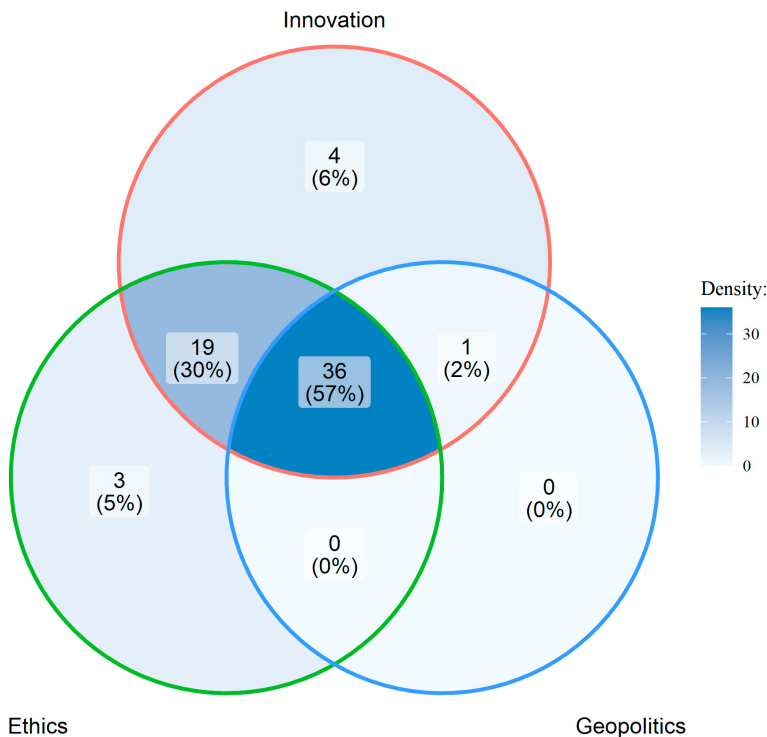


FIGURE 4 Coverage of issue frames in U.S. AI policy documents. Absolute totals with percentages below.

Yet this apparent synthesis risks overstating the relative importance of issue frames other than innovation. When examining the relative amount of attention to each issue frame, allowing for multiple mentions per document, the innovation frame continues to dominate. Treating the dominant frame in a document as the frame mentioned more than all other frames, innovation is dominant in 40 of 63 documents (63.5%), ethics in 12 of 63 documents (19.0%), and geopolitics in only 4 of 63 documents (6.3%).

Importantly, a key related finding is that ethics discourse, while strikingly common, is often subsumed into innovation-oriented goals. This ability to strategically merge and subsume the ethics frame in this way may result from AI's complexity and ambiguity as a general purpose technology with numerous implications, leading to epistemic uncertainty and subsequent interpretive flexibility (Goyal et al., 2021). There are several ways in which this pattern of subsumption is reflected. Ethical issue frames are often featured heavily in introductions, mission statements, and motivations, with less attention in more detailed portions of documents (e.g., policy solutions). Relatedly, ethical issues are often treated as more exploratory, something to be “understood” or concerned with generally. Ethics topics, even when they show up formally in the strategic priority lists of documents, are often further down the list compared to other goals. Finally, even when ethical issues are treated seriously, they are often rendered in the ultimate service of other goals, and thus treated as secondary goals or as means to other ends.

For example, the Executive Order on Maintaining American Leadership in AI (June 2019, p. 20) “emphasizes that maintaining American leadership in AI requires a concerted effort to promote advancements in technology and innovation, *while* protecting civil liberties, privacy, and American values” (emphasis added). The American AI Initiative (September 2019, p. 5) echoes that “in all of these actions, the Initiative emphasizes the importance of advancing AI innovation, *while* fostering public trust and confidence in AI technologies.” Relatedly, ethics is even seen as a barrier to innovation. As the Department of Commerce (DOC) notes (August 2020, p. 1) with respect to explainable AI, “suspicions that the system is biased or unfair ... may slow societal acceptance and adoption of the technology” while the G20 notes similarly (June 2019, p. 5) that securing the trust of the uncertain public “is essential for enabling the benefits of the global digital economy.” Numerous statements echo these examples.

Yet it is important to note that such a conclusion remain provisional as the agenda continues to develop, and that numerous documents and agencies do put forward strong statements surrounding ethics, sometimes with concrete action. Interestingly, some branches of the military are particularly strong in this respect. The Department of Homeland Security (DHS) and Department of Defense (DOD) (February 2020) and Office of the Director of National Intelligence (ODNI) (February and July 2020) have been unusually active in articulating ethical concerns, and the Air Force (September 2019, p. 2) even conveys a precautionary tone to innovation, noting that “artificial intelligence is not the solution to every problem” and that “its adoption must be thoughtfully considered in accordance with our ethical, moral, and legal obligations to the Nation.” Other agencies like the National Highway Traffic Safety Administration (NHTSA) seem to have constructively synthesized goals around ethics (namely, safety, privacy, and access) with goals surrounding innovation. Ultimately however, the evidence supports the concern that prominent discussion of AI ethics remains substantially rhetorical, and that this discourse is being layered into the traditional paradigm of technology governance rather than transforming it.

Policy problems: Traditional emphasis followed by ‘hybrid’ challenges

The next results show the analysis of policy problems in U.S. AI policy documents. The analysis includes only policy problems that received an average of more than one quotation per document.¹⁷ The remaining top 16 problems presented in Figure 5 are classified as traditional, transformative or ‘hybrid.’ This latter category reflects an interesting complexity in AI policy discourse. Namely, while certain topics might be formally associated with AI ethics discourse, such as trust or transparency, the ways in which these topics are discussed often highlight innovation or economic adoption concerns rather than human-centered social and ethical concerns.

For example, ‘trust’ is often rendered as a means to promote innovation, as exemplified previously. Along similar lines, the DHS (December 2020, p. 14) notes that building trust in the American public can “protect and guard against reputational impacts to the Department” while the DOC (August 2020, p. 4) notes that some forms of transparency are “*designed* to generate trust and acceptance by society,” (emphasis added) which according to the DHS (p. 16) should “result in an educated and supportive American public.” Other topics with an ostensibly ethical focus, like privacy and safety, are likewise rarely discussed in terms of human impacts, and are often framed in terms of strategic geopolitical or economic goals. Even topics like diverse participation can be rendered as, on the one hand, inclusion of vulnerable subgroups of the public in decision-making, or alternatively, the need for a diverse array of *experts* across different government agencies and industry bodies. Thus in analyzing topics like privacy, transparency, safety, and diverse participation, each topic is treated as on-balance closest to ‘hybrid’ in its relationship to traditional versus transformative innovation policy.

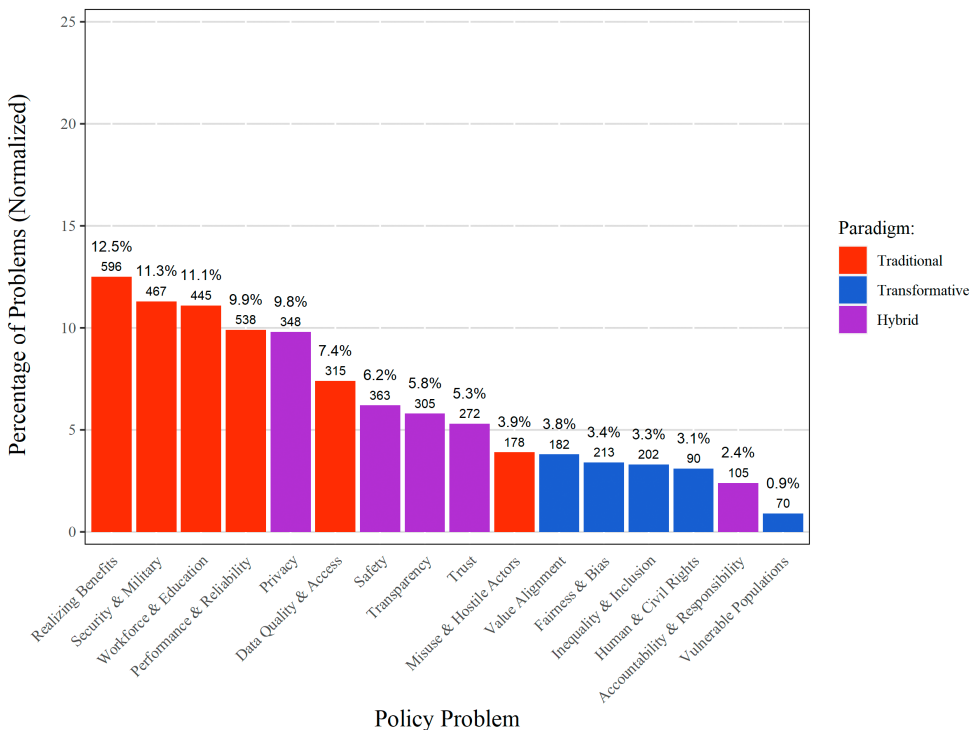


FIGURE 5 Coverage of policy problems in U.S. AI policy documents. Percentages of all policy problems (total $n = 4689$) in documents. Normalized percentages with absolute totals below.

Based on this classification approach, the top several problems are quite traditional in nature. Interestingly, the most common policy ‘problem’ overall is the need to realize the benefits of AI, in light of the opportunity cost of not doing so. This interesting sentiment reflects a strong innovation motive and potentially unique hallmark of technology agenda-setting discourse with respect to problem identification. Meanwhile, problems surrounding the expertise or capacity of the workforce and the need to build better AI systems or attend to security concerns are classically traditional. Overall, a full 44% of problems mentioned fall into the top four categories displayed in Figure 5. In contrast, the transformative problems receive the least attention, such as those surrounding human rights and vulnerable populations. Jointly, fairness and bias, inequality and inclusion, value alignment, human and civil rights, and vulnerable populations constitute only around 16% of total problems discussed in the U.S. AI policy agenda.

Viewed differently, however, social and ethical problems do show up to a surprising degree. For example, while 49 documents mention workforce and education, and 54 mention security or military, as many as 25 documents mention human or civil rights, 28 mention inequality or inclusion, and 41 mention fairness and bias. Yet the distinction between the mere *coverage* by a document of an ethical problem and the *degree* of attention overall is meaningful, a distinction perceivable due to the methodology applied here. One interpretation is that, while ethical problems show up in many documents, they are often treated casually as part of a long list of problems without as detailed attention or consideration. Examining solutions offered in response to these problems is another way of determining the seriousness of policymakers with respect to these problems.

Policy solutions: Traditional solutions with hybrid possibilities

The findings with respect to policy solutions echo the patterns above but suggest new possibilities as well. The top four coded policy solutions, jointly reflecting 53% of all solutions coded, are calls for cooperation (overwhelmingly government and industry), research and development towards increased AI adoption, increasing access to data, and building the STEM and AI workforce. Notably, some of these solutions follow quite linearly from the associated problems: a deficit of high skilled workers implies a need for more training, and the same is true of access to high-quality data (Figure 6).

However, these traditional solutions are followed by a set of practices classified as hybrid here. Numerous documents call for efforts to evaluate the impacts and risks of AI systems, to promote standards and best practices, to increase transparency, and so on. Many of these solutions are theoretically and sometimes explicitly responsive to both traditional and transformative type problems. Thus while calls for formal social and ethical consideration in AI development or public engagement as policy solutions are relatively sparse, there are potential inroads to address transformative type policy concerns. To the extent to which impact assessments, standards, reporting practices, and so on, can be designed so as to foster attention to broader kinds of social and ethical goals, the door remains open to serious action on these dimensions.

Roles of the public and experts: Strongly expert dominated

Finally, Figure 7 reports results related to the role of the public versus that of experts in the U.S. AI policy agenda. The Sankey diagram depicts the overlap between each stakeholder and the 16 policy solutions presented previously. Because the co-occurrence of stakeholders and policy

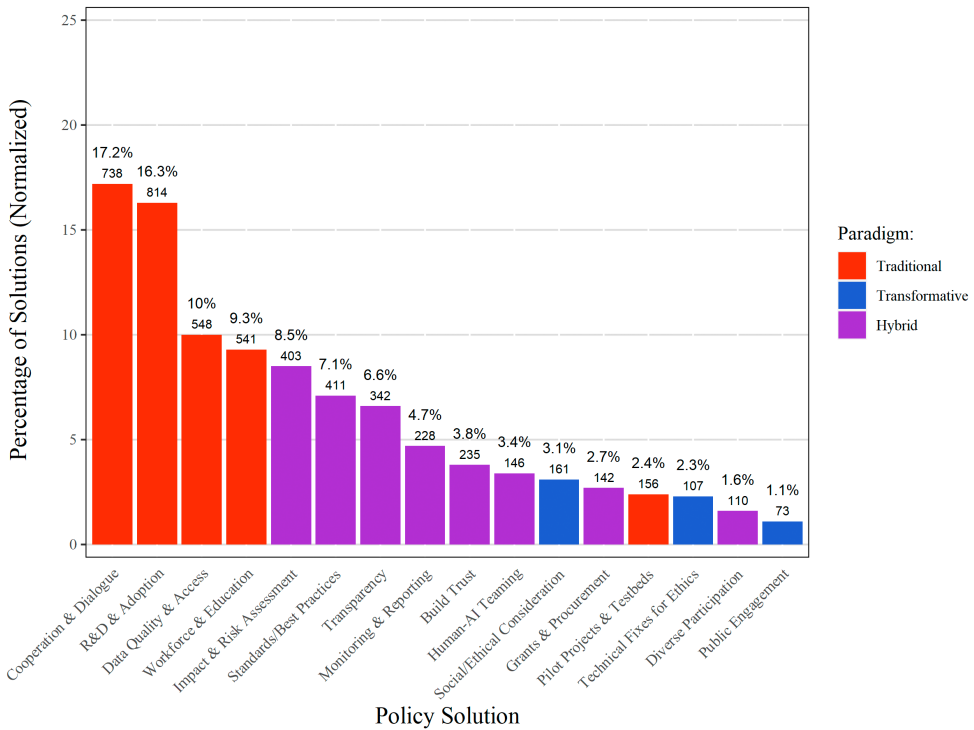


FIGURE 6 Coverage of policy solutions in U.S. AI policy documents. Percentages of all policy solutions (total $n = 5155$) in documents. Normalized percentages with absolute totals below.

solutions was coded at the level of individual paragraphs, this provides only a rough proxy. Paragraphs often involve long list of problems and solutions, and solutions and stakeholders may be mentioned in different sections of documents. Nevertheless, Figure 7 provides a reasonable approximation, and if anything, overstates the role of the public.

In particular, the public receives relatively less attention and takes on a very modest role overall. The public is mentioned in 46 of the 63 documents, while experts are mentioned in 57 documents. Meanwhile, concrete calls for public engagement as a policy solution appear in only 24 documents, and public opinion is mentioned as a factor of interest in only 13 documents. These findings reflect that, while mentions of public good or public involvement show up quite commonly, these calls for action are rarely realized through concrete policy solutions. Members of the public are most often referenced with respect to general dialogue, and even much of this communication is generally one directional: Policymakers discuss taking action to serve the public, to build the public's trust, to educate the public about AI, to foster public adoption of AI, and so on, but offer few concrete plans for two-sided public engagement.

To a lesser extent, members of the public are thought helpful in providing input, at least in theory. Yet few metrics of public opinion are ever presented, suggesting that document writers have spent little time gathering information from public opinion polls, for example. Further, the idea of the 'public' is often construed more broadly than imagined in the transformative paradigm, e.g., with a significant emphasis on members of industry who might comment on regulatory proposals. In contrast, proposals to involve experts from government and industry in AI policy often involve detailed specifications of actors, roles, institutions, objectives, and even timelines. The discrepancy is extremely sharp. This may be explained by beliefs that "industry and

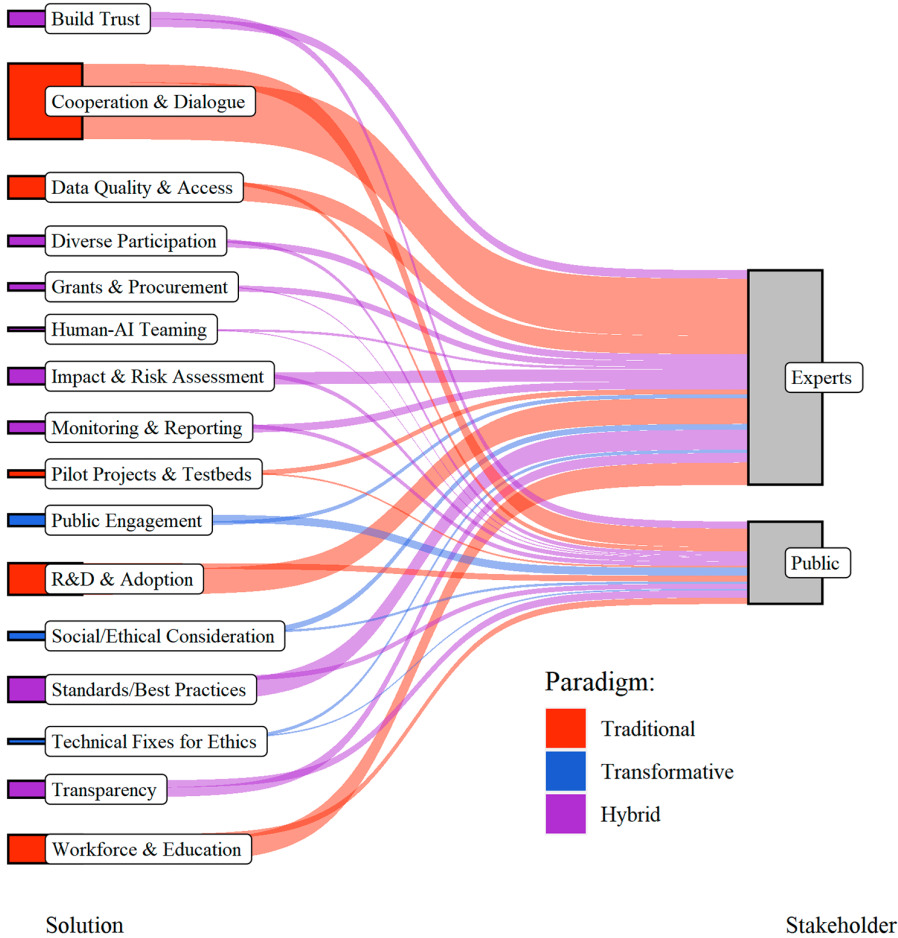


FIGURE 7 Role of the public and experts in AI policy solutions.

academia are the primary sources for emerging AI technologies,” (NSTC, October 2016, p. 34), and that “developing technical expertise will provide the basis for” advances in AI (NSTC, October 2019, p. 39). In combination, it is quite clear that efforts to foster public participation in AI policy (Buhmann & Fieseler, 2022; Crockett et al., 2021) are not being heeded in practice.

DISCUSSION

This article sought to provide theoretical and empirical insight into the development of the U.S. AI policy agenda in light of agenda-setting theory and competition between alternative paradigms of technology governance. However, there are key limitations to this study worth considering. The time period under examination is limited to 2016–2020, whereas many important developments are occurring beyond this time window, including new legislation, strategic documents, discourse around geopolitical competition, and possible transformative ethics-related reforms advancing through agencies like the National Institute of Standards and Technology (NIST) and the Federal Trade Commission (FTC). Moreover, the complexity, nuance, and variation across documents requires some level of interpretation and inference, performed by a single researcher. Other researchers may identify different emphases or ways to aggregate and parse topics of interest, and

may find value in performing additional comparisons across document types, policy subsystems, and over time. Finally, the analysis is largely limited to executive agency documents at the federal level, with some external evidence invoked in a supporting fashion, and with limited coverage of the many actors, institutions, and dynamics involved in agenda-setting. Future research should probe sources like hearings, speeches, legislative text, and documents produced by the public, private, and non-governmental sectors to add insight on the agenda-setting process, as well as employ longitudinal techniques, interviews, and other complementary methods.

Yet reviewing the focusing events, indicators, policy problems and solutions, issue frames, and role of stakeholders in these key strategic AI documents helps to paint an overall picture. Initially, the evidence suggests a striking level of attention to ethics discourse in the U.S. AI policy agenda. Despite the reputation of the United States as comparatively skewed toward innovation compared to, for example, the European Union, numerous government agencies articulate ethical priorities and goals and several even identify their own ethics principles and frameworks. This focus is especially evident in some of the most ‘important’ AI policy documents (identified in Table S1.1)—those with a broad horizontal scope affecting many sectors of government.¹⁸ Developments like the U.S. AI Bill of Rights, the explicit classification of ‘ethics documents’ by the White House, and participation by the United States in ethical statements made by the OECD and G20 are all signs of genuine interest and openness. Yet those invested in the transformative dimensions of AI are likely to be disappointed as these high-level calls for ethics are underserved when documents move beyond broad mission statements, a concern raised in prior literature (Schiff, Rakova, et al., 2021; Ulnicane et al., 2021).

The MSF provides at least two possible theoretical explanations for lack of full translation of ethics into the policy agenda, suggested by the evidence and interpretive work performed here: lack of value acceptability and lack of technical feasibility. Regarding value acceptability, it may be the case that some calls for ethical action go beyond the tolerance of policymakers. For example, more narrowly tailored ethical solutions surrounding issues like privacy or algorithmic bias may be admissible within the bounds of current U.S. values toward technology. Meanwhile, more sweeping calls (Waelen, 2022) to address structural inequality, reform societal power dynamics, and strictly regulate AI from a precautionary perspective may be too outside the bounds for the current U.S. regulatory mode. For instance, while calls for upskilling the workforce are pervasive across the AI agenda, a vanishingly small number of documents call for reforms to the social safety net or redistributive tax reforms.¹⁹ The results here thus provide empirical evidence to support concerns by other scholars (Erman & Furendal, 2022; Wong et al., 2022) that only a subset of AI ethics issues are currently tolerated and translated into practice.

The second potential explanation offered by the MSF and supported here surrounds technical feasibility. Simply, some AI ethics issues are thought to be addressable through technical fixes. It is imagined that new technical practices and standards will help developers avoid unfair, opaque, and privacy-violating AI systems. Indeed, many of the frameworks promoted to address AI ethics (Morley et al., 2021) center around these technical solutions, while broader socio-technical solutions and calls for economic and social reform are scoped out of attention. This sentiment with respect to technical feasibility is expressed clearly by the DOC and NIST (August 2019, p. 16): “While stakeholders ... expressed broad agreement that societal and ethical considerations must factor into AI standards, it is not clear how that should be done and whether there is yet sufficient scientific and technical basis to develop those standards provisions.”

A third and related possible explanation is the lack of appropriate venues to facilitate more expansive transformative goals in light of institutional norms and constraints (Justo-Hanani, 2022). This is evidenced in how calls for ethical action—especially prominent in government-wide AI policy

documents—become narrowed when translated to individual government agencies. For example, addressing vehicle safety and cybersecurity related to autonomous vehicles is comfortable in light of the NHTSA's typical policy rationales and instruments, but addressing widespread societal inequality is not. The same is true of military entities, who are concerned with issues like the trust of autonomous weapons operators and civilian casualties. To the extent to which individual agencies interpret their ethical charges within the bounds of traditional institutional constraints, this seems to delimit the potential for broader transformative change. Appendix S2 provides additional analysis and evidence with respect to the comparison of high-priority government-wide against sector-specific documents, which tentatively supports this explanation.

Importantly, there are various other possible explanations related to the agenda-setting literature, innovation literature, and elsewhere. For example, the state of the agenda could reflect a lack of sufficient or skilled policy entrepreneurship, an inability to achieve coupling, or a balance of political interests favoring the traditional approach to technology policy. These and other explanations require future research to explore.

Another surprising finding is the prevalence of various 'hybrid' problems and solutions, where topics like privacy and safety become 'rote' engineering and software practices, perhaps even justified with respect to strategic economic and geopolitical goals. Yet that these hybrid proposals often become rhetorically divorced from human-centered impacts may ironically be a sign of progress or synthesis (constructive or otherwise). That is, as abstract ethical goals are operationalized into concrete technical proposals surrounding problems like privacy, this act of sensemaking and reduction from rhetoric to action may reflect a necessary normalization of transformative concerns into concrete ways of working (af Malmborg, 2023). Thus, it is not entirely clear if the transformative agenda is being fully—or destructively—subsumed into the traditional paradigm. The emerging agenda may ultimately reflect some degree of layering and productive synthesis.

These considerations point to areas for future action. Stakeholders concerned with AI's social and ethical implications seem to have had substantial success in promoting an associated issue frame, and even strategically merging it with other issue frames. To move from rhetoric to action, however, AI policy entrepreneurs will likely need to find ways to operationalize and translate their problems into workable solutions, especially within policy sectors, and then socialize these ideas to policymakers in relevant government agencies. This could mean working to ensure that emerging proposals surrounding dialogue, public engagement, impact assessment, and standards incorporate social-ethical concerns and strategies. Indeed, the attention of AI standards like IEEE 7010 (Schiff et al., 2020) and NIST's AI Risk Management Framework to socio-technical dimensions of AI seem to be indicators of willingness to push beyond traditional technology governance, though these 'public goods' focused standards efforts also face special challenges (von Ingersleben, 2023). Yet for larger scale structural and social reforms where technical fixes seem unlikely, the barriers to transformation may be even greater. Attempting to deeply change the relationship of the United States to innovation may require long-term socialization, replacement of decision-makers over time, or even an unprecedented crisis that resurfaces socio-ethical concerns.

CONCLUSION

This study contributes to the literature in several ways. To scholars of the policy process, it constitutes one of very few studies addressing technology policy or AI policy in particular, despite the importance of these domains. Methodologically, through both qualitative and quantitative

content analysis, it demonstrates how policy documents can be fruitfully analyzed and interpreted to reveal subtle features of the agenda-setting process. While the large majority of work on the MSF offers qualitative insight based on process tracing and document analysis, this study builds on these methods with extensive mixed methods analysis, showing how the absolute and relative prevalence of elements like focusing events and issue frames can be incorporated into studies of agenda-setting.

For scholars of technology governance and innovation policy, the findings help to answer whether the 21st century is unfolding with respect to a new governance paradigm, as well as how ethics relates to theories of innovation policy. Overall, the evidence for such a grand shift is limited. The rhetorical attention and even willingness to think more broadly about technology's mixed implications suggest a meaningful degree of openness to change, but constraints like technical feasibility and value acceptability in the U.S. context demarcate modern limits. An initial reading of the evidence is that elements of social and ethical concern have been layered into, or are beginning to be synthesized into, the traditional paradigm. Meanwhile, other elements of the transformative agenda, like increased public engagement, have a long way to go.

For scholars and stakeholders of AI policy in particular, the findings also begin to answer key questions like the extent to which ethical concerns in AI are translated into policy, and how the United States may function as an international policy actor. In light of the importance of the United States as a global AI policy leader, the findings offer some notes of promise and some notes of concern. Dismissing the United States as beholden to a neoliberal governance paradigm is too hasty. Indeed, many of the key policy entrepreneurs in the AI ethics community and ongoing work identifying associated policy problems and solutions is occurring in the United States. Moreover, there is (bipartisan) openness to the kinds of socio-ethical perspectives centered in the European Union, reflecting a greater degree of convergence than imagined. How AI governance will unfold in the next decades will thus depend on ongoing efforts to carefully frame problems, develop concrete and feasible solutions, and to identify novel venues and strategies for placing more atypical and transformative strategies on the global policy agenda.

ACKNOWLEDGMENTS

I would like to thank Jason Borenstein, Ellen Zegura, Juan Rogers, Scott Ganz, Natália Bueno, and participants of the workshop on AI Politics and Policy for their helpful feedback on this project.

CONFLICT OF INTEREST

None.

ORCID

Daniel S. Schiff  <https://orcid.org/0000-0002-4376-7303>

ENDNOTES

- ¹ Most typically, research in science and technology using the MSF focuses on the environmental domain (Huber-Stearns et al., 2019). However, also see Goyal et al. (2021) for a recent broader review of the role of technology in the MSF and Justo-Hanani (2022) for a policy process oriented study of AI policy.

- ² This paper defines ‘AI ethics’ with respect to prominent concepts (or ethical topics) that show up commonly across AI ethics frameworks, principles, and policy documents. For example, Schiff, Borenstein, et al. (2021) discuss 25 AI ethics concepts across 112 documents, ranging from fairness and transparency to inequality, human rights, public participation, and vulnerable populations. In other reviews, Fjeld et al. (2020) and Jobin et al. (2019) similarly highlight concepts like responsibility, trust, justice, and transparency, amongst others, as foundational concepts in AI ethics discourse. These concepts form the baseline for the paper’s treatment of how AI ethics manifests or not in U.S. AI policy documents.
- ³ See Appendix S3 for a discussion of the approach to selecting a subset of MSF elements, limitations, and related implications for theory-building.
- ⁴ For example, the OMB’s Memorandum on Guidance for AI Regulation (2020, p. 3) places public participation as its second pillar, arguing that “that public participation ... will improve agency accountability and regulatory outcomes, as well as increase public trust and confidence” and that “agencies must provide ample opportunities for the public to provide information and participate in all stages of the rulemaking process.”
- ⁵ For those interested in a deeper discussion of the study’s use of paradigms and how study findings may influence our understanding, please see Appendix S5.
- ⁶ For accessibility, documents in the dataset referenced in the paper are identified by authoring organization and publication year and month. Readers can cross-reference Tables S1.1 and S1.2 in the Appendix to find the original sources.
- ⁷ Budgetary documents were the least consistent with the rest of the document types in their approach and structure.
- ⁸ The excluded documents are the DOT/FAA Strategic Plan for FY2019-2022 (2018), the G7 Science and Technology Ministers’ Declaration on COVID-19 (2020), and the near duplicate version of the NITRD Artificial Intelligence and Cybersecurity workshop report (2020).
- ⁹ Only a subset of the domains coded are discussed in this paper. See Appendix S3 for a discussion.
- ¹⁰ After coding an initial sample of documents, I recoded a clean subset of documents and compared results to improve coding reliability, important because this is a single-coder effort. Nevertheless, the choice and parsing of codes is at least partially subjective and other researchers may have other preferences. For further information about the addition, removal, and merging of codes, contact the author.
- ¹¹ Of note, the coding process generally excluded appendices, footnotes, generic introductory text (e.g., author or organization biographies), and in cases, sections focused on topics other than AI (e.g., for a document emphasizing AI and quantum computing, it excluded analysis of the latter section).
- ¹² Feedback from experts in policy process theory and AI policy was helpful in improving the codebook however.
- ¹³ For example, a document about bias that mentions bias dozens of times will not exert undue weight.
- ¹⁴ Again, these classifications are only rough heuristics, as individual instances are invariably more complex.
- ¹⁵ The normalization process involves reweighing each document to have the same number of elements within a domain (e.g., focusing events or indicators or issue frames). The total number of elements over all documents is then determined along with the percent attention to each respective element. Note that presenting results based on an absolute or normalized measures produces largely similar results across all analyses.
- ¹⁶ For example, a prototypical statement reflective of which indicators and focusing events are commonly emphasized comes from the Office of the Director of National Intelligence (January 2019, p. iv): “Artificial intelligence (AI) ... has shown dramatic advances in autonomous systems, computer vision, natural language processing, and game playing. These AI systems can perform tasks significantly beyond what was possible only recently (e.g., autonomous systems) and in some cases even beyond what humans can achieve (e.g., chess and Go).”
- ¹⁷ For example, problems such as artificial general intelligence, concerns about monopolies, risks surrounding misinformation, and organizational barriers in AI adoption are excluded with less than 63 mentions each.
- ¹⁸ For example, high-priority government-wide documents include the National AI R&D Plan (October 2016), the Executive Order on Trustworthy Use of AI (December 2020), and the OMB Guidance for Agencies on Regulation of AI (January 2020).
- ¹⁹ These solutions were so sparse that they were eliminated from the analysis of policy solutions.

REFERENCES

- af Malmborg, F. (2023). Narrative dynamics in European Commission AI policy—Sensemaking, agency construction, and anchoring. *Review of Policy Research*.
- Arnstein, S. R. (1969). A ladder of citizen participation. *Journal of the American Institute of Planners*, 35(4), 216–224. <https://doi.org/10.1080/01944366908977225>
- Atkinson, R. D., & Mayo, M. J. (2010). *Refueling the U.S. innovation economy: Fresh approaches to science, technology, engineering and mathematics (STEM) education* (SSRN Scholarly Paper ID 1722822). Social Science Research Network. <https://papers.ssrn.com/abstract=1722822>
- Avenier, M.-J., & Thomas, C. (2015). Finding one's way around various methodological guidelines for doing rigorous case studies: A comparison of four epistemological frameworks. *Systemes d'information Management*, 20(1), 61–98. <https://doi.org/10.3917/sim.151.0061>
- Baumgartner, F. R. (2014). Ideas, paradigms and confusions. *Journal of European Public Policy*, 21(3), 475–480. <https://doi.org/10.1080/13501763.2013.876180>
- Bhaskar, R. (1979). *The possibility of naturalism: A philosophical critique of the contemporary human sciences* (1st ed.). Harvester Press.
- Birkland, T. A. (1998). Focusing events, mobilization, and agenda setting. *Journal of Public Policy*, 18(1), 53–74.
- Birkland, T. A., & Warnement, M. K. (2016). Refining the idea of focusing events in the multiple-streams framework. In R. Zohlnhöfer & F. W. Rüb (Eds.), *Decision-making under ambiguity and time constraints: Assessing the multiple-streams framework* (pp. 91–108). ECPR Press.
- Birks, M., Chapman, Y., & Francis, K. (2008). Memoing in qualitative research: Probing data and processes. *Journal of Research in Nursing*, 13(1), 68–75. <https://doi.org/10.1177/1744987107081254>
- Bouncken, R. B., Qiu, Y., & García, F. J. S. (2021). Flexible pattern matching approach: Suggestions for augmenting theory evolution. *Technological Forecasting and Social Change*, 167, 120685. <https://doi.org/10.1016/j.techfore.2021.120685>
- Buhmann, A., & Fieseler, C. (2022). Deep learning meets deep democracy: Deliberative governance and responsible innovation in artificial intelligence. *Business Ethics Quarterly*, 1–34. <https://doi.org/10.1017/beq.2021.42>
- Cairney, P., & Jones, M. D. (2016). Kingdon's multiple streams approach: What is the empirical impact of this universal theory? *Policy Studies Journal*, 44(1), 37–58. <https://doi.org/10.1111/psj.12111>
- Chuan, C.-H., Tsai, W.-H. S., & Cho, S. Y. (2019). Framing artificial intelligence in American newspapers. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 339–344. <https://doi.org/10.1145/3306618.3314285>
- CIFAR. (2020). *Building an AI world: Report on national and regional AI strategies* (2nd ed., p. 48). CIFAR. <https://cifar.ca/wp-content/uploads/2020/10/building-an-ai-world-second-edition.pdf>
- Cohen, M. D., March, J. G., & Olsen, J. P. (1972). A garbage can model of organizational choice. *Administrative Science Quarterly*, 17(1), 1–25. <https://doi.org/10.2307/2392088>
- Crockett, K., Colyer, E., & Latham, A. (2021). The ethical landscape of data and artificial intelligence: Citizen perspectives. *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–9. <https://doi.org/10.1109/SSCI50451.2021.9660153>
- Dahl, R. A. (1978). Pluralism revisited. *Comparative Politics*, 10(2), 191–203. <https://doi.org/10.2307/421645>
- Daigneault, P.-M. (2014). Reassessing the concept of policy paradigm: Aligning ontology and methodology in policy studies. *Journal of European Public Policy*, 21(3), 453–469. <https://doi.org/10.1080/13501763.2013.834071>
- 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>
- DeLeo, R. A., Taylor, K., Crow, D. A., & Birkland, T. A. (2021). During disaster: Refining the concept of focusing events to better explain long-duration crises. *International Review of Public Policy*, 3(1), Article 1. <https://doi.org/10.4000/irpp.1868>
- 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>
- 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>
- Erkkilä, T. (2023). Global indicators and AI policy: Metrics, policy scripts and narratives. *Review of Policy Research*.
- Erman, E., & Furendal, M. (2022). The global governance of artificial intelligence: Some normative concerns. *Moral Philosophy and Politics*, 9, 267–291. <https://doi.org/10.1515/mopp-2020-0046>

- 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* (p. 82). Centre for the Governance of AI, Future of Humanity Institute, University of Oxford. <https://www.fhi.ox.ac.uk/wp-content/uploads/2021/03/AI-Policy-Levers-A-Review-of-the-U.S.-Governments-tools-to-shape-AI-research-development-and-deployment-%E2%80%93-Fischer-et-al.pdf>
- 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. <https://papers.ssrn.com/abstract=3518482>
- Fletcher, A. J. (2017). Applying critical realism in qualitative research: Methodology meets method. *International Journal of Social Research Methodology*, 20(2), 181–194. <https://doi.org/10.1080/13645579.2016.1144401>
- Freeman, R., & Maybin, J. (2011). Documents, practices and policy. *Evidence & Policy: A Journal of Research, Debate and Practice*, 7(2), 155–170. <https://doi.org/10.1332/174426411X579207>
- Galanos, V. (2019). Exploring expanding expertise: Artificial intelligence as an existential threat and the role of prestigious commentators, 2014–2018. *Technology Analysis & Strategic Management*, 31(4), 421–432. <https://doi.org/10.1080/09537325.2018.1518521>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15–31. <https://doi.org/10.1177/1094428112452151>
- Goyal, N., Howlett, M., & Taihagh, A. (2021). Why and how does the regulation of emerging technologies occur? Explaining the adoption of the EU general data protection regulation using the multiple streams framework. *Regulation & Governance*, 15(4), 1020–1034. <https://doi.org/10.1111/rego.12387>
- Greer, S. (2016). John W. Kingdon, agendas, alternatives, and public policies. In M. Lodge, E. C. Page, & S. J. Balla (Eds.), *The Oxford handbook of classics in public policy and administration* (Vol. 1). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199646135.013.18>
- Grillitsch, M., Hansen, T., & Madsen, S. (2018). *How novel is transformative innovation policy?* (No. 2020/08; Papers in Innovation Studies, p. 18). Centre for Innovation, Research and Competence in the Learning Economy (CIRCLE).
- Hall, P. A. (1993). Policy paradigms, social learning, and the state: The case of economic policymaking in Britain. *Comparative Politics*, 25(3), 275–296. <https://doi.org/10.2307/422246>
- Herweg, N. (2016). Clarifying the concept of policy-communities in the multiple-streams framework. In R. Zohlnhöfer & F. W. Rüb (Eds.), *Decision-making under ambiguity and time constraints: Assessing the multiple-streams framework* (pp. 125–145). ECPR Press.
- Herweg, N., Huß, C., & Zohlnhöfer, R. (2015). Straightening the three streams: Theorising extensions of the multiple streams framework. *European Journal of Political Research*, 54(3), 435–449. <https://doi.org/10.1111/1475-6765.12089>
- Herweg, N., Zahariadis, N., & Zohlnhöfer, R. (2017). The multiple streams framework: Foundations, refinements, and empirical applications. In C. M. Weible & P. A. Sabatier (Eds.), *Theories of the policy process* (4th ed., pp. 17–53). Westview Press.
- Hoddy, E. T. (2019). Critical realism in empirical research: Employing techniques from grounded theory methodology. *International Journal of Social Research Methodology*, 22(1), 111–124. <https://doi.org/10.1080/13645579.2018.1503400>
- Hsieh, H.-F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative Health Research*, 15(9), 1277–1288. <https://doi.org/10.1177/1049732305276687>
- Huber-Stearns, H. R., Schultz, C., & Cheng, A. S. (2019). A multiple streams analysis of institutional innovation in forest watershed governance. *Review of Policy Research*, 36(6), 781–804. <https://doi.org/10.1111/ropr.12359>
- Imbrie, A., Gelles, R., Dunham, J., & Aiken, C. (2021). *Contending frames: Evaluating rhetorical dynamics in AI*. Center for Security and Emerging Technology. <https://doi.org/10.51593/20210010>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Jones, M. D., Peterson, H. L., Pierce, J. J., Herweg, N., Bernal, A., Raney, H. L., & Zahariadis, N. (2016). A river runs through it: A multiple streams meta-review. *Policy Studies Journal*, 44(1), 13–36. <https://doi.org/10.1111/psj.12115>
- Justo-Hanani, R. (2022). The politics of artificial intelligence regulation and governance reform in the European Union. *Policy Sciences*, 55(1), 137–159. <https://doi.org/10.1007/s11077-022-09452-8>

- Kim, J. (2023). Containing artificial intelligence: Traveling AI-essentialism and IT imaginaries in South Korea and France. *Review of Policy Research*.
- Kingdon, J. W. (1984). *Agendas, alternatives, and public policies* (1st ed.). Little, Brown.
- Knaggård, Å. (2015). The multiple streams framework and the problem broker. *European Journal of Political Research*, 54(3), 450–465. <https://doi.org/10.1111/1475-6765.12097>
- Köstler, L., & Ossewaarde, R. (2022). The making of AI society: AI futures frames in German political and media discourses. *AI & Society*, 37(1), 249–263. <https://doi.org/10.1007/s00146-021-01161-9>
- 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>
- Lincoln, Y. S., & Guba, E. G. (1986). But is it rigorous? Trustworthiness and authenticity in naturalistic evaluation. *New Directions for Program Evaluation*, 1986(30), 73–84. <https://doi.org/10.1002/ev.1427>
- 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>
- Mazzucato, M. (2018). Mission-oriented innovation policies: Challenges and opportunities. *Industrial and Corporate Change*, 27(5), 803–815. <https://doi.org/10.1093/icc/dty034>
- McGregor, S. (2020). Preventing repeated real world AI failures by cataloging incidents: The AI Incident Database. *ArXiv:2011.08512 [Cs]*. <http://arxiv.org/abs/2011.08512>
- Miles, M. B., Huberman, A. M., & Saldana, J. (2013). *Qualitative data analysis*. SAGE.
- Miller, H. T. (2015). Introduction to the symposium, part 2: Interrogating neutral public administration. *Administrative Theory & Praxis*, 37(4), 223–226. <https://doi.org/10.1080/10841806.2015.1083821>
- Mintrom, M., & Norman, P. (2009). Policy entrepreneurship and policy change. *Policy Studies Journal*, 37(4), 649–667. <https://doi.org/10.1111/j.1541-0072.2009.00329.x>
- Morlacchi, P., & Martin, B. R. (2009). Emerging challenges for science, technology and innovation policy research: A reflexive overview. *Research Policy*, 38(4), 571–582. <https://doi.org/10.1016/j.respol.2009.01.021>
- Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2021). Operationalising AI ethics: Barriers, enablers and next steps. *AI & Society*. <https://doi.org/10.1007/s00146-021-01308-8>
- National AI Initiative Office, Office of Science and Technology Policy. (2022). *National Artificial Intelligence Initiative documents*. <https://www.ai.gov/documents/>
- Neri, H., & Cozman, F. (2020). The role of experts in the public perception of risk of artificial intelligence. *AI & Society*, 35(3), 663–673. <https://doi.org/10.1007/s00146-019-00924-9>
- OECD. (2021). *State of implementation of the OECD AI principles: Insights from national AI policies* (No. 311; p. 93). OECD. <https://doi.org/10.1787/1cd40c44-en>
- Ouchchy, L., Coin, A., & Dubljević, V. (2020). AI in the headlines: The portrayal of the ethical issues of artificial intelligence in the media. *AI & Society*, 35(4), 927–936. <https://doi.org/10.1007/s00146-020-00965-5>
- Owen, R., Stilgoe, J., Macnaghten, P., Gorman, M., Fisher, E., & Guston, D. (2013). A framework for responsible innovation. In R. Owen, J. Bessant, & M. Heintz (Eds.), *Responsible innovation* (pp. 27–50). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781118551424.ch2>
- Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., Lyons, T., Manyika, J., Mishra, S., & Nibbles, J. C. (2019). *Artificial Intelligence Index 2019 annual report* (p. 291). Human-Centered AI Institute, Stanford University. https://hai.stanford.edu/sites/g/files/sbiybj10986/f/ai_index_2019_report.pdf
- Perry, B., & Uuk, R. (2019). AI governance and the policymaking process: Key considerations for reducing AI risk. *Big Data and Cognitive Computing*, 3(2), Article 2. <https://doi.org/10.3390/bdcc3020026>
- Pierre, J. (Ed.). (2000). *Debating governance: Authority, steering, and democracy*. Oxford University Press.
- Princen, S., & 't Hart, P. (2014). Putting policy paradigms in their place. *Journal of European Public Policy*, 21(3), 470–474. <https://doi.org/10.1080/13501763.2013.876177>
- Roberts, N. C., & King, P. J. (1991). Policy entrepreneurs: Their activity structure and function in the policy process. *Journal of Public Administration Research and Theory: J-PART*, 1(2), 147–175.
- 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. (2022). Education for AI, not AI for education: The role of education and ethics in national AI policy strategies. *International Journal of Artificial Intelligence in Education*, 32(3), 527–563. <https://doi.org/10.1007/s40593-021-00270-2>

- Schiff, D., Ayesh, A., Musikanski, L., & Havens, J. C. (2020). IEEE 7010: A new standard for assessing the well-being implications of artificial intelligence. *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2746–2753. <https://doi.org/10.1109/SMC42975.2020.9283454>
- 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. <https://doi.org/10.1109/TTS.2021.3052127>
- Schiff, D., Rakova, B., Ayesh, A., Fanti, A., & Lennon, M. (2021). Explaining the principles to practices gap in AI. *IEEE Technology and Society Magazine*, 40(2), 81–94. <https://doi.org/10.1109/MTS.2021.3056286>
- Schot, J., & Steinmueller, W. E. (2018). Three frames for innovation policy: R&D, systems of innovation and transformative change. *Research Policy*, 47(9), 1554–1567. <https://doi.org/10.1016/j.respol.2018.08.011>
- Sinkovics, N. (2018). Pattern matching in qualitative analysis. In C. Cassell, A. Cunliffe, & G. Grandy (Eds.), *The SAGE handbook of qualitative business and management research methods: Methods and challenges* (pp. 468–484). SAGE Publications Ltd. <https://doi.org/10.4135/9781526430236.n28>
- 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>
- 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. https://doi.org/10.1007/978-3-030-56286-1_9
- 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>
- Stirling, A. (2016). Addressing scarcities in responsible innovation. *Journal of Responsible Innovation*, 3(3), 274–281. <https://doi.org/10.1080/23299460.2016.1258946>
- Stivers, C. (2015). Rule by nobody: Bureaucratic neutrality as secular theodicy. *Administrative Theory & Praxis*, 37(4), 242–251. <https://doi.org/10.1080/10841806.2015.1083825>
- Stix, C. (2021). Actionable principles for artificial intelligence policy: Three pathways. *Science and Engineering Ethics*, 27(1), 15. <https://doi.org/10.1007/s11948-020-00277-3>
- Stix, C., & Maas, M. (2021). Bridging the gap: The case for an ‘incompletely theorized agreement’ on AI policy. *AI and Ethics*, 1, 261–271. <https://doi.org/10.1007/s43681-020-00037-w>
- Taeiigh, A. (2021). Governance of artificial intelligence. *Policy and Society*, 40(2), 137–157. <https://doi.org/10.1080/14494035.2021.1928377>
- Tarnoff, B., Redwine, C., Tan, J., Sheets, K., Nedzhvetskaya, N., & Rajasekaran, S. (2021). Collective action in tech. *Archive—Collective Action in Tech*. <https://data.collectiveaction.tech/>
- Ulicane, 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>
- Ulicane, 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>
- Ulicane, 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(2), 158–177. <https://doi.org/10.1080/14494035.2020.1855800>
- Ulicane, I., Knight, W., Leach, T., Stahl, B. C., & Wanjiku, W.-G. (2022). Governance of artificial intelligence: Emerging international trends and policy frames. In M. Tinnirello (Ed.), *The global politics of artificial intelligence* (pp. 29–55). Chapman and Hall/CRC.
- United States Innovation and Competition Act of 2021, S. 1260, 117th Cong., Senate, 117. (2021). (testimony of Charles E. Schumer). <https://www.congress.gov/bill/117th-congress/senate-bill/1260>
- 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>
- von Ingersleben, N. (2023). Competition and cooperation in artificial intelligence standard-setting: Explaining emerging patterns. *Review of Policy Research*.

- Waelen, R. (2022). Why AI ethics is a critical theory. *Philosophy & Technology*, 35(1), 9. <https://doi.org/10.1007/s13347-022-00507-5>
- Warnke, P., Koschatzky, 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 No. 49). Fraunhofer ISI Discussion Papers—Innovation Systems and Policy Analysis. <https://www.econstor.eu/handle/10419/129191>
- White, M. D., & Marsh, E. E. (2006). Content analysis: A flexible methodology. *Library Trends*, 55(1), 22–45. <https://doi.org/10.1353/lib.2006.0053>
- Wong, R. Y., Madaio, M. A., & Merrill, N. (2022). Seeing like a toolkit: How toolkits envision the work of AI ethics. *ArXiv:2202.08792 [Cs]*. <http://arxiv.org/abs/2202.08792>
- Wynn, D., & Williams, C. K. (2012). Principles for conducting critical realist case study research in information systems. *MIS Quarterly*, 36(3), 787–810. <https://doi.org/10.2307/41703481>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE.
- Zhang, D., Maslej, N., Brynjolfsson, E., Etchemendy, J., Lyons, T., Manyika, J., Ngo, H., Niebles, J. C., Sellitto, M., Sakhaee, E., Shoham, Y., Clark, J., & Perrault, R. (2022). *Artificial intelligence index report 2022* (p. 230). AI Index Steering Committee, Human-Centered AI Institute, Stanford University. https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf

AUTHOR BIOGRAPHY

Daniel S. Schiff is an Assistant Professor of Technology Policy at Purdue University's Department of Political Science. At Purdue, he is co-director of GRAIL, the Governance and Responsible AI Lab. As a policy scientist with a background in philosophy, Daniel's studies the formal and informal governance of AI through policy and industry, as well as AI's social and ethical implications in domains like education, manufacturing, finance, and criminal justice. His research addresses topics such as industry standards and organizational practices for AI ethics, public and elite opinion and influence dynamics in the policy process, the development of social responsibility attitudes amongst future computing and engineering professionals, and the role of the public in governing emerging technologies. Previously, Daniel served as Responsible AI Lead at JP Morgan Chase & Co and as Director of Research, Evaluation, and Planning at the Philadelphia Education Fund. He studied Philosophy at Princeton University, focusing on robotics and intelligent systems, before completing an MS in Social Policy at the University of Pennsylvania, and a PhD in Public Policy from the Georgia Institute of Technology. He has published in venues such as *Public Administration*, *Technology in Society*, *AI & Society*, the *International Journal of AI in Education*, *IEEE Transactions on Technology & Society*, the *Journal of Engineering Education*, the *AMA Journal of Medical Ethics*, and *Nanotechnology* and can be reached at dschiff@purdue.edu or [@dan_schiff](https://twitter.com/dan_schiff).

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Schiff, D. S. (2023). Looking through a policy window with tinted glasses: Setting the agenda for U.S. AI policy. *Review of Policy Research*, 00, 1–28. <https://doi.org/10.1111/ropr.12535>