

Bridging AI Research and Policy: A Co-Production Framework and the Case for Policy Transfer Offices

Malihe Alikhani¹, Jeongeun Kim^{*2}, Anna Haensch^{*3}, Saki Imai^{*1},
Serena Booth⁴, Andrea Stith¹, Adriana Bankston⁵

¹Northeastern University

²University of Maryland, College Park

³University of Wisconsin - Madison

⁴Brown University

⁵STEM Advocacy Institute
m.alikhani@northeastern.edu

Abstract

Artificial intelligence systems now shape governance, labor, healthcare, and public life at scale, yet there remains a persistent gap between technical knowledge and the policy processes meant to govern AI's use. We argue that this gap is not a communication failure but an *institutional design problem*: the absence of sustained, equitable mechanisms that make researcher–policymaker collaboration feasible. Drawing on theories of science–policy co-production, we present empirical evidence from a survey of AI researchers (n=44) that identifies the dominant structural barriers as time costs (66%), pathway opacity (57%), and weak institutional support, not lack of access to policymakers. We introduce a co-production framework centered on the novel proposal for *Policy Transfer Offices* (PTOs): institutional structures, analogous to technology transfer offices, that broker relationships, provide translation templates, and legitimize policy work within academic evaluation systems. We further expand on the proposal, in the dimensions of placement, governance, funding, and failure modes, and ground it in two in-depth case studies (algorithmic collusion and AI in education) that illustrate when technical expertise fails to reach governance in time, and what institutionalized co-production enables. We further analyze equity in participation, showing how current ad hoc models concentrate policy influence among senior researchers at elite institutions, and propose reforms for broadening access. The paper offers empirical evidence, an actionable institutional mechanism, and equity-centered design principles to align AI research with democratic governance.

1 Introduction

Artificial intelligence systems are embedded in domains governed by public policy such as housing, education, energy, labor, healthcare, and public services. Governments around the world are moving to legislate, regulate, and operationalize AI governance. Despite a growing volume of AI research relevant to these decisions, a persistent gap remains between technical knowledge and the policy processes that govern its use.

^{*}These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This gap reflects deep structural misalignments between how AI research is produced, evaluated, and incentivized, and how policy decisions are actually made. A growing body of work on fairness, accountability, and transparency has documented the societal risks of AI systems, from algorithmic bias and opacity to market concentration and democratic erosion (Jobin, Ienca, and Vayena 2019; Raji et al. 2020; Costanza-Chock, Raji, and Buolamwini 2022). At the same time, policy scholarship has shown that evidence does not flow linearly from research to legislation (Head 2015; Cairney 2016). Policymaking is inherently political and iterative, shaped by institutional constraints, timing, and relationships. When AI researchers engage in policy through one-off briefs or reactive interventions, their contributions often arrive too late or lack the contextual fit required for durable implementation and effectiveness.

The dominant barrier to effective AI governance, we argue, is not a lack of research, but the absence of institutionalized mechanisms that make sustained, equitable, and usable engagement possible. Effective governance requires what Jasanoff 2004 calls *co-production*: an iterative collaboration where researchers and policymakers jointly shape problem definitions and evidence standards over time, rather than treating policy as a downstream application of technical work.

To ground this claim, we present findings from a survey of AI researchers examining their experiences with policy engagement. The results reveal a striking pattern: while the majority view policy impact as important (51%), their engagements are constrained by time costs (66%), pathway opacity (57%), and limited institutional support (32%), not notably, by lack of access. These constraints systematically privilege senior researchers at elite institutions, narrowing whose expertise informs policy and undermining equity in AI governance (Ban, Park, and You 2023; Clauset, Arbesman, and Larremore 2015).

Building on these findings, we propose a co-production framework that shifts responsibility for engagement from individuals to institutions. Central to this is the novel proposal for *Policy Transfer Offices* (PTOs), institutional structures that reduce transaction costs, broker relationships between

researchers and policymakers, provide translation templates, and create mechanisms for recognizing policy engagement within academic careers. Unlike fellowships that temporarily relocate individuals or government relations offices that primarily represent an institution's official interests, PTOs embed durable infrastructure within research institutions to support the translation of diverse scholarly expertise into policy.

We illustrate how co-production succeeds or fails through two in-depth case studies: algorithmic collusion in antitrust law, and AI governance in education. These cases demonstrate that when technical expertise is integrated early and iteratively, governance outcomes are more responsive and evidence-based; when it is not, algorithmic practices become entrenched before law can intervene.

This paper makes four contributions:

1. **Empirical evidence:** We characterize structural frictions—specifically time costs and pathway opacity—that prevent diverse AI researchers from engaging in policy.
2. **Institutional design:** We advance a concrete, operationalized model for PTOs, including placement, governance, incentives, funding, and failure modes.
3. **Analytic case studies:** We develop two cases that illuminate what institutional mediation enables and where its absence creates harm.
4. **Equity-centered governance:** We show how current models concentrate influence and propose reforms for broader participation.

2 Related Work

Models of Science–Policy Interaction

Early models of science–policy interaction assumed a linear pathway where research findings flow directly to decision-makers, and have been extensively critiqued as oversimplistic (Head 2015; Cairney 2016). Jasanoff's 2004 foundational work on co-production established that scientific knowledge and social order are mutually constitutive: effective governance requires ongoing collaboration rather than unidirectional transfer. Cash et al. 2003 identify credibility, salience, and legitimacy as key attributes that determine whether knowledge systems effectively inform decisions. These same notions can be directly applied to AI policy.

Oliver, Lorenc, and Innvær 2014 conclude that successful translation depends on relationships, timing, and mutual understanding. More recent scholarship emphasizes iterative engagement; Maas, Pauwelussen, and Turnhout 2022 find that linear models persist even when co-production is espoused, with unclear roles perpetuating ineffective knowledge transfer. Scheufele 2014 argues that science communication is inherently political, yet AI researchers often default to the discredited “deficit model” (National Academies of Sciences, Engineering, and Medicine 2017), which treats policymakers as information-deficient rather than as partners with different but legitimate epistemic standards.

AI Governance Challenges

AI governance presents unique obstacles: rapid technological changes that outpace regulation, deep uncertainty about system capabilities, and the need for technical literacy among policymakers (Calo 2017; Bommasani et al. 2021). The proliferation of AI ethics guidelines reveals convergence in principles but divergence in implementation (Jobin, Ienca, and Vayena 2019). Recent regulatory developments—the EU AI Act (European Parliament and Council of the European Union 2024) and the NIST AI Risk Management Framework (National Institute of Standards and Technology 2023)—demonstrate growing appetite for comprehensive governance while highlighting the need for technical expertise in implementation (Novelli et al. 2025).

Work in fairness, accountability, and transparency has documented social risks and proposed governance interventions (Raji et al. 2020; Selbst et al. 2019). Costanza-Chock, Raji, and Buolamwini 2022 find that audit-centered accountability may foreclose more fundamental questions about whether particular systems should exist at all. Taei-hagh 2021; 2025 and Ulnicane and Erkkilä 2023 document the political economy of AI governance, showing how institutional power shapes whose expertise counts in regulatory processes, a topic of direct relevance to our equity analysis. Yang et al. 2024 identify structural barriers to HCI–policy collaboration that mirror our survey findings, including institutional misalignment and the absence of durable bridging mechanisms.

Institutional Mechanisms: Knowledge Brokers and Technology Transfer

Several mechanisms are currently being used to bridge research and policy. Fellowship programs such as the American Association for the Advancement of Sciences (AAAS) Science and Technology Policy Fellowships (American Association for the Advancement of Science 2018) and TechCongress (TechCongress 2024) place researchers in government, while advisory bodies such as the European Centre for Algorithmic Transparency (European Commission 2023) provide ongoing technical expertise for regulatory enforcement. Knowledge brokers and boundary organizations have emerged as critical intermediaries (MacKillop et al. 2023). Gluckman, Bardsley, and Kaiser 2021 emphasize trust, transparency, and the synthesis of pluralistic evidence for effective brokerage. Singh et al. 2014 find that policy engagement is insufficiently recognized by institutions, with time constraints and lack of institutional support representing the most significant barriers. This conclusion is strongly corroborated by our survey.

Technology transfer offices (TTOs) provide a compelling organizational analogy. Established to bridge university research and commercial applications, TTOs reduce transaction costs by centralizing expertise, relationships, and legal infrastructure (Macho-Stadler, Pérez-Castrillo, and Veugelers 2007). We draw on this model to propose PTOs for research–policy translation. Critically, while TTOs themselves have a mixed record—with success rates varying by institutional context and incentive structure (Macho-Stadler,

Pérez-Castrillo, and Veugelers 2007)—the organizational logic of centralizing translation expertise applies strongly to the policy domain, where individual researchers face identical coordination failures.

Equity in Research–Policy Engagement

Access to policy engagement is not equally distributed. The career stage of a researcher is a limiting factor: senior researchers receive more invitations to policy conversations than early-career scholars (Ban, Park, and You 2023). Institutional prestige further stratifies engagement: elite research universities maintain disproportionate policy connections (Clauaset, Arbesman, and Larremore 2015). Geographic proximity to political centers provides additional structural advantages.

Representative bureaucracy scholarship demonstrates that inclusive participation in policy design improves outcomes in efficacy, efficiency, and legitimacy (Cingolani 2023). For AI policy specifically, where disparate impact across demographic groups is a persistent concern (Selbst et al. 2019), ensuring diverse researcher participation is substantively important for governance quality, not merely procedurally desirable. Chambers et al. 2021 analyze how power dynamics shape whose knowledge counts in collaborative research, underscoring the need for explicit equity mechanisms.

What Is Specific to AI?

A fair question is whether the transfer problem we diagnose is specific to AI at all, or simply a general feature of science–policy interfaces. Our answer is that the mechanism is general, but the severity and stakes are distinctive. The most obvious difference is the pace. AI capabilities diffuse into deployed products in months, while legislative and regulatory cycles run for years. Medicine mediates a similar tension through standing evidence-synthesis institutions such as FDA advisory committees and NICE, and climate science does so through assessment bodies such as the IPCC. AI has no equivalent standing infrastructure, which is precisely the gap that PTOs would address.

The second difference is that AI is a general-purpose technology. A drug or an emissions source implicates one regulatory domain; a single foundation model implicates housing, education, labor, and healthcare policy at once, so relevant expertise must reach many uncoordinated policy venues simultaneously. This multiplies the transaction costs that our survey documents. Third, the technical capacity to evaluate frontier systems is disproportionately located within the private firms being governed (Abdalla and Abdalla 2021). Independent academic input is therefore structurally scarcer than in medicine or climate, where public research infrastructure is mature. Finally, the behavior of AI systems deployed is often not fully characterizable even by their developers. Policy-relevant knowledge in this domain is genuinely co-produced through auditing and evaluation, rather than settled upstream and then handed over.

These comparisons also indicate what AI governance can borrow from older interfaces. The IPCC shows that periodic, institutionalized synthesis can discipline a fast-moving evidence base. Health technology assessment shows that trans-

lation offices with standardized formats can make technical evidence legible to non-specialist decision-makers. PTOs adapt both lessons to a domain that currently lacks them. Conversely, if the PTO model proves out, it should be transferred to other emerging-technology domains facing the same institutional vacuum.

3 Survey: Researcher Perspectives on Policy Engagement

While prior work has proposed models for research–policy engagement and documented mechanisms for AI governance, limited empirical evidence characterizes how AI researchers experience structural frictions, incentive misalignments, access constraints, and other practical dynamics of participation.

Methods

We distributed an online survey to AI researchers through academic networks and conference channels between November 2025 and May 2026. The survey included questions about demographic and professional background, prior engagement with policy activities, barriers to engagement, desired institutional support, and preferred communication formats. It was administered via Qualtrics with IRB approval. Participation was voluntary and responses were collected anonymously. We limit our sample to fully completed responses ($n=44$). For analysis, we performed a thematic pass over open-ended responses to identify recurring barriers and illustrative quotes; we do not claim saturation or statistical representativeness given the sample size. The full survey instrument is provided in the supplementary appendix to support interpretation and reuse.

Sample Characteristics

The majority of the respondents (78%) were affiliated with academic institutions, with the remainder distributed across independent research (17%), industry (5%), think tanks (2%), and government (2%). Thirty-three percent had six or more years of experience in AI research. Only 30% had prior experience working for the U.S. federal government. The primary research domains included machine learning (24%), policy and social implications (16%), natural language processing (14%), ethics and fairness (9%), computer vision (7%), and privacy and security (2%). Because our recruitment focused on US and European networks, findings should be interpreted in that context; researchers from the Global South or outside these networks may face a different, potentially more acute set of barriers.

Findings

The respondents rated the importance of influencing policy for their career goals and their confidence in communicating with policymakers on a 5-point scale (Figure 1). The majority rated policy influence as important (4–5; 51%), corresponding to the share that reported high confidence in communication (4–5; 51%). However, policy importance elicited substantial neutrality (37% selected 3). We interpret this




Figure 1: Diverging stacked bar chart comparing policy importance vs communication confidence ratings ($n = 44$). Negative ratings (1–2) are shown on the left, neutral (3) is centered, and positive ratings (4–5) are on the right. Both measures show 51% positive ratings (4–5), but policy importance shows stronger consensus with only 14% rating it low (1–2) compared to 25% lacking communication confidence. The 37% neutral response for policy importance suggests significant uncertainty about its career relevance.




Figure 2: Barriers to policy engagement ($n=44$, multiple selections allowed). Time constraints and pathway opacity dominate; access to policymakers is the least-cited barrier.




Figure 3: Desired institutional support ($n=44$, multiple selections allowed). Networking, recognition in evaluation, and dedicated funding are most requested, which maps directly onto the dominant barriers.

with caution. Neutrality is consistent with a genuine disinterest in policy influence. But read alongside the open-ended responses, in which respondents described wanting to engage without knowing whether it would “count,” and alongside the high demand for evaluation recognition (Finding 2), it is at least partly attributable to uncertainty about whether policy engagement aligns with career incentives. Disentangling these interpretations would require follow-up items that we did not include in the survey, and we flag this as a limitation. Low ratings (1–2) were more common for confidence in communication (25%) than for policy importance (14%), indicating that some researchers care about policy impact but feel unprepared to engage effectively.

Finding 1: The main bottlenecks are time and pathway opacity, not access. Respondents selected from multiple barriers (Figure 2). Time constraints dominated (66%), reflecting competition with research, teaching, and service. More than half reported uncertainty about how to engage (57%), describing difficulty identifying entry points, audiences, or appropriate forms of contribution. Resource limitations (34%) and lack of institutional support (32%) were also common. Notably, access barriers were the least frequently cited (23%), suggesting that the bottleneck is not reaching

policymakers but navigating the path to doing so productively. Open-ended responses reinforced this: one respondent wrote, “*I don’t know the level of pitch/communication they already want/need... I haven’t figured out how to find the right audience.*” Another noted structural advantage: “*Higher-resourced institutions obviously have more avenues to engage in these efforts.*”

Finding 2: Researchers want institutional support that reduces transaction costs and makes engagement legible within careers. The most requested supports were connections and networking (52%), recognition in evaluation (48%), and dedicated funding (48%), followed by templates and examples (45%) and training and workshops (43%) (Figure 3). These requests map directly onto dominant barriers: networking responds to pathway opacity; recognition addresses career legibility; funding supports time costs. The demand for templates further suggests that researchers want concrete, reusable forms for translating technical work into policy-relevant outputs, not just generic communication advice.

Finding 3: Willingness is high but conditional on support. Despite substantial barriers, respondents expressed

strong willingness to engage when engagement is feasible and supported. One respondent wrote, “*I see the work we do as impacting real-world issues,*” while another noted that their research was “*adopted in some contexts*” with policy implications. These responses align with the quantitative pattern that barriers reflect feasibility and institutional support, not lack of motivation.

Implications

Taken together, the survey reveals a gap between interest in policy impact and the ability to act on that interest. The dominant constraints—time costs and unclear engagement pathways—compound the limited institutional infrastructure to shift policy engagement toward researchers with discretionary time, established networks, or institutional privilege. These findings motivate institutional mechanisms that reduce transaction costs and make engagement pathways more navigable and equitable.

An important scope note: this survey captures only the researcher side of the co-production equation. Understanding the full barriers to co-production ultimately requires studying policymakers’ institutional dynamics—how bureaucracies process external expertise, what they need from researchers, and how political timing shapes demand for evidence. We position policymaker-side analysis as critical complementary future work, not a missing variable in our diagnosis.

4 The Co-Production Framework

We define co-production as an institutionalized process in which AI researchers and policymakers jointly shape problem definitions, evidence standards, and policy options through sustained interaction. Moreover, we explore the notion that research informs policy not only through published findings but through structured pathways that support translation, interpretation, and iteration.

Our survey results indicate that limited engagement is not primarily driven by lack of interest or technical competence, but by three recurring constraints: (i) high time costs, (ii) unclear or informal engagement pathways, and (iii) weak alignment between policy engagement and academic evaluation. When these constraints are left unaddressed, engagement depends on personal networks and discretionary resources, privileging senior researchers and well-resourced institutions. A co-production framework must therefore shift responsibility for engagement from individuals to institutions.

Lowering Transaction Costs

For many researchers, policy engagement requires navigating unfamiliar institutions, identifying appropriate policy audiences, and producing non-academic outputs—often without logistical support. These requirements impose transaction costs that compete directly with research and teaching. Institutions can reduce these costs through centralized support: identifying relevant policy venues, coordinating meetings or testimony preparation, and supporting policy-relevant materials. By reducing overhead, institutions make

participation feasible for researchers currently excluded by time or resource constraints.

Making Engagement Pathways Navigable

Survey respondents frequently reported uncertainty about how to engage, even when motivated. Informal pathways favor researchers with prior experience or insider networks, limiting the diversity of expertise in policy discussions. Effective co-production requires clearly defined engagement pathways specifying how researchers can contribute at different stages of the policy process: agenda setting, drafting, implementation, and evaluation. Explicit pathways reduce reliance on personal connections and lower barriers for early-career researchers and those at less-resourced institutions.

Aligning Incentives with Engagement

Policy engagement carries opportunity costs rarely recognized in academic evaluation. When policy-facing work is treated as extracurricular, researchers must absorb these costs individually. Institutions can address this misalignment by recognizing policy engagement as a legitimate scholarly activity and service in a way that is sufficient to ensure that researchers who contribute are not penalized relative to peers focused exclusively on traditional outputs.

Mutual Capacity Building

Co-production is not a one-directional fix. Researchers and policymakers need different but complementary capacities. Researchers need skills in translating technical uncertainty, structuring policy-legible arguments, and understanding regulatory timelines. Policymakers benefit from baseline technical literacy, structured relationships with research communities, and clear channels for research input on pending legislation. Both the survey literature and our cases show that engagement is most durable when both sides develop these capacities together, rather than depending on one side to absorb the full translation burden.

5 Policy Transfer Offices: Design and Implementation

We propose Policy Transfer Offices (PTOs) as a concrete institutional mechanism to operationalize co-production (Figure 4). PTOs are analogous to technology transfer offices (TTOs) that support research commercialization (Macho-Stadler, Pérez-Castrillo, and Veugelers 2007): where TTOs manage intellectual property and industry partnerships, PTOs support the translation and circulation of research into governance contexts.

Critically, we do not claim that PTO-like functions are entirely absent. Fellowship programs, boundary organizations, and advisory bodies already perform some of this work in fragmented form. Our contribution is the claim that *consolidation and incentive alignment* are what existing mechanisms lack: unlike fellowships that temporarily relocate individuals, PTOs embed durable infrastructure within research institutions, making policy engagement systematically available and legible within careers.




Figure 4: Policy Transfer Office as institutional hub. Solid arrows form the institutional pipeline; the dashed curve marks the informal direct channel that currently privileges well-connected researchers. The PTO reduces transaction costs, centralizes translation expertise, and creates accountable pathways for co-production.

Core Functions

PTOs would serve four core functions:

Brokerage. Proactive matchmaking between researchers and legislative staff, regulatory agencies, or executive bodies—bypassing the “access opacity” that our survey found. For early-career scholars without established networks, the PTO functions as an institutional introduction that personal social capital cannot provide.

Translation Support. Templates, editorial reviews, and writing assistance for policy briefs, testimony, and regulatory comments. This addresses the finding that 45% of researchers want reusable templates. Policy-legible formats are genuinely different from formats used for academic outlets, and learning them imposes a time cost that a PTO can amortize across many researchers.

Impact Tracking. Systematic documentation of “legislative citations,” advisory appointments, and regulatory comments informed by affiliated researchers. This serves two purposes: creating a feedback loop so researchers understand impact and providing institutional evidence for recognizing policy work in tenure and promotion.

Equity Mechanisms. Dedicated pathways for early-career researchers, researchers in teaching-focused institutions, and researchers from underrepresented groups. This could include: internal travel and stipend support for policy convenings; a “policy mentorship” matching program pairing early-career scholars with senior researchers experienced in engagement; and rotation programs placing graduate students in policy organizations or legislative offices for short-term.

Institutional Placement

Where should a PTO sit? We identify three viable models, with different tradeoffs:

- **Research Office placement** (e.g., under a VP for Research): highest legitimacy within the institution, easiest integration with grant reporting and IP frameworks, but risk of being siloed from faculty governance.
- **Provost-level placement:** strongest signal that policy engagement is valued institutionally, suitable for large R1 universities with active federal relationships.
- **Consortium or regional model:** multiple institutions pool a single PTO, allowing teaching-focused or smaller institutions to access infrastructure they could not sustain independently. This is our recommended model for broadening equity; it decouples PTO access from institutional prestige.

Two scope conditions bound this design. First, policy-advisory ecosystems vary considerably across countries. Some European systems have comparatively formalized channels for scientific advice, such as parliamentary technology assessment offices, while others, including the United States, lean on ad hoc testimony and interest-group intermediation. PTOs are not a one-size-fits-all structure. Their brokerage function should attach to whatever advisory pathway is dominant in a given jurisdiction, and the consortium model in particular can be adapted to national research-council systems where individual universities lack direct policy relationships.

Second, PTOs are not think tanks nor government relations offices. Think tanks produce policy research and pursue their own institutional agendas, while government relations offices may connect faculty with policymakers but generally represent the university’s official legislative, regulatory, and funding interests. A PTO is an agenda-neutral infrastructure that lowers the cost of engagement for any affiliated researcher, much as a TTO does not invent itself

but enables inventors. The distinction matters for equity: access to think tanks reproduces the same network effects our survey documents, whereas a PTO's mandate is precisely to route around them.

A related objection is that universities already hold the relevant capacities. Research development offices know which faculty produce policy-relevant work and already build “why this matters” narratives; government affairs holds legislative fluency and staffer relationships; technology transfer has practiced skill in distilling technical work for non-expert audiences; sponsored research tracks the federal regulatory landscape and holds relevant compliance expertise. Our claim is not that these assets are absent but that they are dispersed across offices each optimized for a different outcome, funder appeal, institutional interest, licensing revenue, compliance, with no unit accountable for policy translation as such. The open design question is therefore ownership and coordination: whether a PTO is a dedicated office or a coordinating function with an explicit mandate over capacities that already exist. The defining feature is the mandate, not new capability.

Funding and Sustainability

PTOs require dedicated, stable funding. We envision three potential streams: (1) institutional overhead on relevant federal grants (e.g., a small fraction of F&A recoveries directed to the PTO); (2) foundation grants specifically targeting research-policy translation (several major foundations already fund this work in fragmented ways); and (3) federal supplemental mechanisms akin to NSF's “Broader Impacts” requirement, but with institutional-level rather than project-level accountability.

The gap is one of research content as much as infrastructure. Public funders have supported work adjacent to regulatory need, transparency and interpretability research among it, but the quantitative estimates on which regulatory decisions actually turn are seldom the object of funding themselves. What is missing is closer to a science of regulating AI: measurement and estimation designed for the decisions regulators face rather than repurposed from adjacent literatures after the fact. One plausible path to a first PTO is therefore a supplement to an existing translation-infrastructure award of this kind, which could also lend external legitimacy to a function that might otherwise read as purely internal.

Without dedicated funding, policy engagement will remain concentrated among researchers with flexible budgets or institutional privilege—precisely the inequity we seek to address.

Failure Modes and Safeguards

We take seriously that TTOs have a mixed record (Macho-Stadler, Pérez-Castrillo, and Veugelers 2007), and that PTO-like structures could fail or cause damages. We identify three principal failure modes:

Industry capture. A PTO could become a conduit for industry-funded research to gain disproportionate policy access. Safeguards include: transparent disclosure of all funding sources and affiliations; an independent advisory board

with civil society representation; and explicit policies prohibiting the PTO from brokering relationships on behalf of research with undisclosed industry conflicts.

Elite concentration. A PTO housed at a single prestigious institution could replicate the prestige hierarchy it aims to disrupt. The consortium model addresses this structurally; PTOs should also track demographic composition of brokered engagements and report publicly.

Metric distortion. If “legislative citations” become a high-stakes metric in promotion decisions, researchers may optimize for visible or media-facing work at the expense of substantively important but less legible contributions. PTO impact tracking should be used for documentation and feedback, not comparative ranking.

6 Case Studies in Co-Production

We present two in-depth case studies to illustrate when co-production fails and what institutionalized mediation enables. Both cases draw on published research, legal records, and policy documents; they function as analytic illustrations rather than causal evaluations of specific interventions.

Case 1: Algorithmic Collusion and Antitrust Law

The governance problem. Markets are increasingly coordinated by AI and algorithms rather than by human actors, while the law is still written as if collusion takes form through meetings, phone calls, or explicit human-to-human agreements (Booth 2025). The RealPage case provides a prominent illustration: an AI-based pricing system used by large landlords is alleged to have synchronized rent increases and vacancy targets across entire cities (Vogell, Coryne, and Little 2022). In 2024, the U.S. Department of Justice and a coalition of state attorneys general sued RealPage under the Sherman Act (U.S. Department of Justice 2024)—the 1890 statute whose core premise is that competition requires prohibiting *coordinated conduct*. While the case has since settled (U.S. Department of Justice 2025), it functions as a case study of how AI price-setting systems can reshape competitive dynamics in ways the law was not designed to detect.

The technical-legal mismatch. The Sherman Act is keyed to the wrong technical boundary. Traditional antitrust doctrine focuses on agreements, communications, and explicit information sharing, because it was written to target coordination among human actors. Modern machine learning systems do not require any of these to produce coordinated outcomes. With access to rich market data, a sufficiently capable model can infer demand, occupancy, and competitor behavior well enough to sustain supra-competitive prices without any non-public information exchange. This makes a structurally new form of coordination possible: firms can each deploy independent systems, claim to be acting autonomously, and still reach a stable high-price equilibrium through learning dynamics and feedback effects. Experimental and theoretical work already suggests such tacit, algorithm-mediated collusion is plausible (Arunachaleswaran et al. 2024; Den Boer, Meylahn, and Schinkel 2022; Bernhardt and Dewenter 2020).

The timing problem. Researchers were producing relevant work on algorithmic pricing and collusion before the RealPage conduct was alleged (Vogell, Coryne, and Little 2022; Den Boer, Meylahn, and Schinkel 2022). But this research was not integrated into regulatory frameworks proactively: enforcement moved reactively, after the AI system had been deployed at scale across a major housing market. The enforcement challenge—how to establish coordinated conduct when each firm’s AI acts “independently”—reflects the absence of legal standards developed with technical input before deployment.

What co-production would have enabled. A PTO-brokered relationship between antitrust economists, ML researchers, and DOJ staff during the *design* of algorithmic pricing standards could have: (1) contributed technical definitions of “coordination” applicable to ML systems; (2) established transparency and auditing requirements for AI pricing tools as a condition of market use; and (3) developed evidentiary frameworks for detecting tacit collusion that do not require proving explicit agreement. The antitrust scholars and ML researchers who eventually informed the RealPage coalition did so through informal, ad hoc channels—precisely the pathway our survey found most respondents cannot access. Institutionalized brokerage would have made this expertise available earlier and more broadly, including to state-level enforcers without access to elite networks.

What this case shows. Algorithmic antitrust is not yet a domain where institutional co-production has succeeded. It illustrates the cost of reactive engagement: by the time courts intervene, algorithmic practices may be deeply embedded across an industry. The relevant research on equilibrium dynamics, auditing methodologies, and evidentiary standards exists. The institutional pathway to make it actionable does not.

Case 2: AI and Education

A contrasting institutional model. Education governance provides a striking contrast to antitrust. Rather than relying on courts and ex post enforcement, AI governance in education has largely unfolded through guidance, task forces, and intermediary institutions (Schiff 2022; Ifenthaler et al. 2024). This policy style creates more opportunities for iterative engagement, making it a useful case for examining how co-production can be partially institutionalized and where its limits lie.

Research integration in practice. State and national AI task forces, agency working groups, and cross-sector committees have served as venues where researchers, educators, and policymakers jointly interpret evidence and define guardrails. State-level AI task forces documented by the Education Commission of the States explicitly draw on research to prioritize privacy protections, equity, and educator capacity building. Reports from the Georgia Senate Study Committee on Artificial Intelligence (Georgia Senate Study Committee on Artificial Intelligence 2024) and the Arkansas AI and Analytics Center of Excellence (Office of the Governor of Arkansas 2025) similarly emphasize phased adoption, pilot testing, and ongoing evaluation over rapid deployment. California’s guidance on AI in educa-

tion reflects research-informed concerns around data governance and bias mitigation (California Department of Education 2025). At the national level, intermediary organizations translate research into policy-relevant guidance (Darling-Hammond 2025; American Institutes for Research 2025).

What works and what doesn’t. The education case demonstrates that standing institutions can lower the costs of research–policy engagement: task forces absorb uncertainty, update guidance over time, and incorporate evolving research without requiring individual researchers to navigate policy processes on their own. Core policy concerns include data privacy, algorithmic trustworthiness, and equity (Ifenthaler et al. 2024; Holmes et al. 2022).

However, the education case also reveals the limits of ad hoc co-production. The rapid public release of generative AI tools produced a governance crisis that existing task force structures were not fast enough to address. Many educators report that institutional guidance arrived post-hoc, after students were already using LLMs in ways that outpaced policy. The AFT–Microsoft partnership controversy illustrates the risk of industry capture in intermediary institutions: when boundary organizations are funded by the same industry whose products they help govern, their legitimacy is compromised.

What this case shows. Education illustrates that intermediary structures can make research input more continuous than episodic—a genuine achievement relative to the antitrust case. But co-production requires active design of conflict-of-interest safeguards, and task forces cannot substitute for the durable institutional infrastructure that a PTO would provide: a permanent brokerage function, not a temporary committee.

Cross-Case Lessons

Across both cases, the recurring constraint is not the availability of relevant research but whether institutions provide durable pathways for researchers and policymakers to iterate on problem definitions, evidence standards, and implementable options. Antitrust illustrates the cost when pathways are absent. Education illustrates partial success and the design requirements—transparency, conflict-of-interest safeguards, sustained rather than episodic engagement—that distinguish effective co-production from co-optation.

7 Equity and Access

Effective AI policy requires equitable participation in its development. Currently, opportunities for policy engagement are not equally distributed.

Current Inequities

Career stage is one limiting factor. Senior researchers often have more access and are invited into policy conversations more often than early-career scholars (Ban, Park, and You 2023). Institutional prestige further concentrates engagement: elite research universities have more policy connections than teaching-focused or under-resourced institutions (Clauzet, Arbesman, and Larremore 2015). From 1960 to 2018, 8.45% of invited congressional witnesses came from

universities or think tanks (Ban, Park, and You 2023), and it is reasonable to expect that policymakers preferentially call on researchers from elite institutions due to name recognition. Geographic proximity to political centers provides additional structural advantages.

This concentration of influence is not merely a logistical byproduct but a sociotechnical gatekeeping mechanism. By relying on informal, elite networks, the current policy–science interface enforces a prestige bias that prioritizes institutional status over technical diversity. The result is what we call *distilled expertise*: the unique risks faced by marginalized communities, often best understood by researchers outside elite circles, are filtered out before they reach legislative text.

The asymmetry is starkest when academic engagement is set against private-sector influence. Technology firms run substantial, professionalized lobbying operations. Major AI developers and their trade associations spend tens of millions of dollars a year on U.S. federal lobbying, keep permanent government-affairs staff, and benefit from direct relationships between executives and senior officials (Popiel 2018; Khanal, Zhang, and Taeihagh 2025; Patnaik et al. 2026). Industry also shapes the research base itself through funding relationships with academic AI labs (Abdalla and Abdalla 2021). An individual researcher facing the time costs and pathway opacity our survey documents cannot match that kind of sustained presence through occasional effort, however motivated. The result is not just underrepresentation of academic expertise but a systematic tilt in whose evidence reaches policymakers: targeted research and framing from parties with commercial stakes in the regulatory outcome. This asymmetry is itself a central motivation for PTOs. Durable, institutionally supported academic engagement is the counterweight that episodic individual outreach cannot provide. Without it, the evidentiary environment of AI policymaking defaults to the best-funded voices.

Why Diverse Participation Matters

Inclusive participation in AI policy design matters substantively, not merely procedurally. Representative bureaucracy scholarship demonstrates that diverse participation improves policy efficacy, efficiency, and legitimacy (Cingolani 2023). For AI policy specifically, where disparate impact across demographic groups is a persistent concern (Selbst et al. 2019), the researcher community informing governance should reflect the diversity of communities affected by AI systems. The stakes of this filtering are concrete. Demographic disparities in facial recognition accuracy were documented by researchers outside the developing firms (Buolamwini and Gebru 2018), and those systems have since been implicated in wrongful arrest of a Black man in the United States (Hill 2020). Automated eligibility and fraud-detection systems in public benefits have imposed severe harms on low-income families. The Dutch childcare benefits scandal is the most prominent example: an algorithmic risk model disproportionately flagged families of immigrant background, and the fallout contributed to the government’s resignation (Amnesty International 2021). These are harms of a kind scholars studying poverty governance had long warned

about (Eubanks 2018). In each case, the affected communities and the researchers closest to them held the relevant knowledge earliest, and in each case that knowledge reached policy only after deployment and injury. Early-career researchers bring perspectives shaped by recent technical developments, and researchers at teaching institutions often have deeper ties to the communities most affected by AI deployment. Broadening participation is therefore a mechanism for surfacing these risks before entrenchment, not merely a procedural virtue.

Recommendations for Equity

We propose three categories of reform:

Funding mechanisms: Travel support, stipends, and time buyouts for researchers from underrepresented groups and teaching-focused institutions; grant supplements specifically for policy translation activities.

Pipeline programs: Fellowships and training designed for researchers without existing policy connections, with explicit outreach to early-career scholars and underrepresented groups.

Institutional mandates: Guidelines for advisory bodies requiring demographic and institutional diversity in membership; PTO reporting on the composition and nature of brokered policy engagements.

Equity in Policy Outcomes

Beyond participation, equity concerns extend to policy outcomes. Evaluating equity in AI policy outcomes is highly domain-dependent. In healthcare, notions of equity are tied to access and delivery as mediated by social determinants (Osonuga et al. 2025; Gurevich, El Hassan, and El Morr 2023). Such analyses require not only substantial high-quality data, but also technologists, domain experts, and quantitative tools—precisely the multi-stakeholder co-production our framework supports. Ensuring AI policy benefits are distributed fairly across communities requires cooperation from multiple stakeholders, often including federal agencies like the NIH and statistical agencies. Building accountability mechanisms accessible to affected communities—not just researchers and policymakers—is essential.

8 Discussion

This paper advances a claim about why AI research so often fails to shape legislation. The dominant barrier is not a lack of researcher interest or technical capability, but a lack of institutional capacity that makes engagement feasible, repeatable, and broadly accessible. In our survey, interest in policy impact is high, and self-reported confidence is comparable. The binding constraints are limited time, unclear pathways, and weak institutional support that forces policy work to compete with publishing, teaching, and service.

The survey results clarify what “support” must mean in practice. Respondents did not primarily describe policymakers as unreachable. Instead, they requested mechanisms that reduce overhead and make engagement legible within careers: introductions and networking, recognition in evalu-

ation, dedicated funding, and reusable templates. The implication is clear: pressures for researchers to do more outreach will not change participation patterns unless institutions lower the costs of engagement and clarify how researchers can contribute at different stages of the policy process.

Our results and case studies suggest several elements of effective co-production. First, **sustained relationships matter**: one-time engagements are less effective than ongoing partnerships because policy problems, regulatory timelines, and implementation details evolve over months or years, and technical input is most usable when it can be revisited and refined. Second, **co-production requires mutual capacity building**: researchers and policymakers need different but complementary skills, including translation across epistemic standards, uncertainty communication, and operationalization of technical evidence into actionable options. Third, **institutional support is not optional**: the same transaction costs respondents identified (time, pathway opacity, lack of support) will predictably reproduce unequal participation unless they are reduced and redistributed. Fourth, **incentives must be aligned**: if policy engagement does not count in academic careers, participation remains contingent on discretionary time and informal prestige-based invitations. Fifth, **co-production must be designed for equitable access**: opportunities should extend beyond elite institutions and senior researchers, both because legitimacy depends on broad participation and because diverse expertise is necessary to anticipate disparate impacts.

Limitations

This work has limitations that bound our claims. The survey sample ($n=44$) is modest and recruited primarily from US and European networks; findings may not generalize to AI researchers in other contexts, and self-selection bias may overrepresent researchers already interested in policy. The case studies draw on published research, legal records, and policy documents rather than on interviews or ethnographic observation; they function as analytic illustrations, not causal evaluations. The PTO framework requires empirical validation in real institutional settings, and some recommendations may face implementation barriers not fully addressed here, such as institutional inertia and funding scarcity.

These limitations point to future work that is concrete rather than aspirational. The natural next step is a designed pilot: implementing PTO models at two or three institutions of deliberately different types, including at least one R1 university and one teaching-focused institution participating through the consortium model, over a three-year horizon with success criteria specified in advance. Candidate metrics include the number and career-stage composition of brokered engagements relative to the ad hoc baseline our survey characterizes, researcher time-to-first-engagement, and the institutional diversity of researchers reaching policy venues. A pilot of this kind would turn the framework advanced here from an argument into a testable intervention. Complementary work includes developing validated approaches to documenting policy impact without leaning on simplistic

metrics, extending the analysis to global contexts where advisory ecosystems differ substantially, and studying policy-maker perspectives directly, through interviews with legislative staff and agency officials, to understand what forms of evidence, timing, and relationships actually shape demand for research input.

9 Conclusion

This paper argues that the research-policy gap in AI is an institutional design problem. When engagement depends on individual outreach, technical expertise arrives too late, in forms difficult to use in policy settings, from a narrow slice of the research community. We advance a co-production framework that treats policy engagement as shared institutional infrastructure rather than individualized outreach. The framework shifts responsibility from researchers to institutions that can lower costs, make pathways navigable, and ensure that policy-facing work is legible within academic careers.

We operationalize this shift through the proposal of Policy Transfer Offices: institutional structures that broker relationships, provide translation support, track impact, and create mechanisms for recognition and resourcing. We develop this proposal with enough specificity to be actionable: including institutional placement options, funding mechanisms, equity safeguards, and failure modes to avoid. The case studies illustrate why this matters: across domains, the recurring constraint is whether institutions enable researchers and policymakers to iterate on problem definitions, evidence, and implementation options *before* systems become entrenched. When these pathways are absent, policy engagement privileges researchers with time, prestige, and proximity to power, excluding perspectives necessary to anticipate disparate impacts and govern AI systems equitably.

The structures built now will shape whose knowledge counts and how the benefits of AI are distributed. A co-production infrastructure grounded in equity is not an addition to responsible AI governance; it is a precondition for it.

Ethical Statement

This research was conducted with IRB approval. Survey participation was voluntary and anonymous; no personally identifiable information was collected. The survey was distributed through academic networks and conference channels, and findings are interpreted as exploratory given sample size and recruitment method. We acknowledge that our positionality as researchers embedded in academic institutions shapes how we diagnose barriers and design interventions; our proposed solutions reflect an institutional vantage point that may underweight perspectives from community organizations, government practitioners, or industry researchers. We have sought to address this through an explicit equity analysis and by positioning community feedback as a core governance input, but we recognize this is a starting point rather than a resolution.

References

- Abdalla, M.; and Abdalla, M. 2021. The Grey Hoodie Project: Big Tobacco, Big Tech, and the Threat on Academic Integrity. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 287–297. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384735.
- American Association for the Advancement of Science. 2018. Science & Technology Policy Fellowships.
- American Institutes for Research. 2025. Developing Principles to Guide Artificial Intelligence (AI) Use in Education Research. Technical report, American Institutes for Research.
- Amnesty International. 2021. Xenophobic Machines: Discrimination through Unregulated Use of Algorithms in the Dutch Childcare Benefits Scandal. Technical Report EUR 35/4686/2021, Amnesty International.
- Arunachaleswaran, E. R.; Collina, N.; Kannan, S.; Roth, A.; and Ziani, J. 2024. Algorithmic collusion without threats. *arXiv preprint arXiv:2409.03956*.
- Ban, P.; Park, J. Y.; and You, H. Y. 2023. How are politicians informed? Witnesses and information provision in congress. *American Political Science Review*, 117(1): 122–139.
- Bernhardt, L.; and Dewenter, R. 2020. Collusion by code or algorithmic collusion? When pricing algorithms take over. *European Competition Journal*, 16(2–3): 312–342.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; et al. 2021. On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258*.
- Booth, S. 2025. Position: Strong Consumer Protection is an Inalienable Defense for AI Safety in the United States. In *Forty-second International Conference on Machine Learning Position Paper Track*.
- Buolamwini, J.; and Gebru, T. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 77–91. PMLR.
- Cairney, P. 2016. *The Politics of Evidence-Based Policy Making*. Springer.
- California Department of Education. 2025. Artificial Intelligence — Guidance for the Safe and Effective Use of AI in California Public Schools. Technical report.
- Calo, R. 2017. Artificial Intelligence Policy: A Primer and Roadmap.
- Cash, D. W.; Clark, W. C.; Alcock, F.; Dickson, N. M.; Eckley, N.; Guston, D. H.; Jäger, J.; and Mitchell, R. B. 2003. Knowledge systems for sustainable development. *Proceedings of the National Academy of Sciences*, 100(14): 8086–8091.
- Chambers, J. M.; Wyborn, C.; Ryan, M. E.; Reid, R. S.; Riechers, M.; Serban, A.; Bennett, N. J.; Cvitanovic, C.; Fernández-Jiménez, M. E.; Galvin, K. A.; et al. 2021. Six modes of co-production for sustainability. *Nature Sustainability*, 4(11): 983–996.
- Cingolani, L. 2023. Representative bureaucracy and perceptions of social exclusion in Europe: Evidence from 27 countries. *Administration & Society*, 55(3): 515–540.
- Clauset, A.; Arbesman, S.; and Larremore, D. B. 2015. Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1): e1400005.
- Costanza-Chock, S.; Raji, I. D.; and Buolamwini, J. 2022. Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1571–1583.
- Darling-Hammond, L. 2025. *Educating in the AI Era: The Urgent Need to Redesign Schools*.
- Den Boer, A. V.; Meylahn, J. M.; and Schinkel, M. P. 2022. Artificial collusion: Examining supracompetitive pricing by Q-learning algorithms. Technical report, Tinbergen Institute Discussion Paper.
- Eubanks, V. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. USA: St. Martin's Press, Inc. ISBN 1250074312.
- European Commission. 2023. About ECAT: European Centre for Algorithmic Transparency.
- European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act).
- Georgia Senate Study Committee on Artificial Intelligence. 2024. Final Report of the Senate Study Committee on Artificial Intelligence (SR 476). Technical report, Georgia General Assembly.
- Gluckman, P. D.; Bardsley, A.; and Kaiser, M. 2021. Brokerage at the science–policy interface: from conceptual framework to practical guidance. *Humanities and Social Sciences Communications*, 8(1): 84.
- Gurevich, E.; El Hassan, B.; and El Morr, C. 2023. Equity within AI systems: What can health leaders expect? In *Healthcare Management Forum*, volume 36, 119–124. SAGE Publications.
- Head, B. W. 2015. Policy Analysis: Evidence Based Policy-Making. In Wright, J. D., ed., *International Encyclopedia of the Social & Behavioral Sciences*, 281–287. Elsevier.
- Hill, K. 2020. Wrongfully Accused by an Algorithm. *The New York Times*.
- Holmes, W.; Porayska-Pomsta, K.; Holstein, K.; Sutherland, E.; Baker, T.; Shum, S. B.; Santos, O. C.; Rodrigo, M. M. T.; Cukurova, M.; Bittencourt, I. I.; et al. 2022. Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32(3): 504–526.
- Ifenthaler, D.; Majumdar, R.; Gorissen, P.; Judge, M.; Mishra, S.; Raffaghelli, J.; and Shimada, A. 2024. Artificial intelligence in education: Implications for policymakers, researchers, and practitioners. *Technology, Knowledge and Learning*, 29(4): 1693–1710.
- Jasanoff, S. 2004. *States of Knowledge: The Co-Production of Science and the Social Order*. Routledge.

- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9): 389–399.
- Khanal, S.; Zhang, H.; and Taeihagh, A. 2025. Why and how is the power of Big Tech increasing in the policy process? The case of generative AI. *Policy and Society*, 44(1): 52–69.
- Maas, T. Y.; Pauwelussen, A.; and Turnhout, E. 2022. Co-producing the science–policy interface: towards common but differentiated responsibilities. *Humanities and Social Sciences Communications*, 9(1): 1–11.
- Macho-Stadler, I.; Pérez-Castrillo, D.; and Veugelers, R. 2007. Licensing of university inventions: The role of a technology transfer office. *International Journal of Industrial Organization*, 25(3): 483–510.
- MacKillop, E.; Connell, J.; Downe, J.; and Durrant, H. 2023. Making sense of knowledge-brokering organisations: boundary organisations or policy entrepreneurs? *Science and Public Policy*, 50(6): 950–960.
- National Academies of Sciences, Engineering, and Medicine. 2017. *Communicating Science Effectively: A Research Agenda*. Washington, DC: National Academies Press.
- National Institute of Standards and Technology. 2023. Artificial Intelligence Risk Management Framework (AI RMF 1.0).
- Novelli, C.; Hacker, P.; Morley, J.; Trondal, J.; and Floridi, L. 2025. A Robust Governance for the AI Act: AI Office, AI Board, Scientific Panel, and National Authorities. *European Journal of Risk Regulation*, 16(2): 566–590.
- Office of the Governor of Arkansas. 2025. Arkansas AI and Analytics Center of Excellence Initial Report. Technical report.
- Oliver, K.; Lorenc, T.; and Innvær, S. 2014. New directions in evidence-based policy research: a critical analysis of the literature. *Health Research Policy and Systems*, 12(1): 34.
- Osonuga, A.; Osonuga, A. A.; Fidelis, S. C.; Osonuga, G. C.; Juckes, J.; and Olawade, D. B. 2025. Bridging the digital divide: artificial intelligence as a catalyst for health equity in primary care settings. *International Journal of Medical Informatics*, 106051.
- Patnaik, A.; Sullivan, B.; Sunderland, I.; and Beckel, M. 2026. Lobbying Disclosures Reveal Big Tech Spends More Than \$230,000 per Day to Buy Influence.
- Popiel, P. 2018. The Tech Lobby: Tracing the Contours of New Media Elite Lobbying Power. *Communication, Culture and Critique*, 11(4): 566–585.
- Raji, I. D.; Smart, A.; White, R. N.; Mitchell, M.; Gebru, T.; Hutchinson, B.; Smith-Loud, J.; Theron, D.; and Barnes, P. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44.
- Scheufele, D. A. 2014. Science Communication as Political Communication. *Proceedings of the National Academy of Sciences*, 111(Supplement 4): 13585–13592.
- 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.
- Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68.
- Singh, G. G.; Tam, J.; Sisk, T. D.; Klain, S. C.; Mach, M. E.; Martone, R. G.; and Chan, K. M. A. 2014. A more social science: barriers and incentives for scientists engaging in policy. *Frontiers in Ecology and the Environment*, 12(3): 161–166.
- Taeihagh, A. 2021. Governance of artificial intelligence. *Policy and Society*, 40(2): 137–157.
- Taeihagh, A. 2025. Governance of Generative AI. *Policy and Society*, 44(1): 1–22.
- TechCongress. 2024. TechCongress Congressional Innovation Fellowship.
- Ulnicane, I.; and Erkkilä, T. 2023. Politics and policy of Artificial Intelligence. *Review of Policy Research*, 40(5): 612–625.
- U.S. Department of Justice. 2024. United States of America et al. v. RealPage Inc. Case No. 1:24-cv-00710.
- U.S. Department of Justice. 2025. Justice Department Requires RealPage to End Sharing of Competitively Sensitive Information.
- Vogell, H.; Coryne, H.; and Little, R. 2022. Rent Going Up? One Company’s Algorithm Could Be Why. ProPublica.
- Yang, Q.; Wong, R. Y.; Jackson, S. J.; Junginger, S.; Hagan, M. D.; Gilbert, T.; and Zimmerman, J. 2024. The future of HCI-policy collaboration. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–15.