Stanford Law School
Law and Policy Lab

AI in Criminal Justice:

Why Governance Matters and

How to Make it Work

A report by Stanford Law & Policy Lab

"Governing the Use of Artificial Intelligence in Criminal Justice" (Law 810C Fall 2025)

March 2026

https://law.stanford.edu/education/only-at-sls/law-policy-lab/
559 Nathan Abbot Way Stanford, CA

Authors

Isabel Astrachan (JD ’27), Tejaswita Kharel (LLM ’26), Archit Lohani (LLM ’26), Anna Makena McGuire (JD ’27), Alexander Rivkin (JD ’27), Emi Sakamoto (BA ’28), Bianka Sedmakova (Visiting Exchange Student from University of Vienna ’26) , Mirac Suzgun (JD / Ph.D. Computer Science ’26) , Victor Y. Wu (JD ‘25 / Ph.D. Political Science ’28)

Instructors: Professors Jonathan J. Wroblewski and Robert Weisberg

Supporting Administrator: Luciana Herman, Program Director, Law & Policy Lab

Acknowledgment

This report is the work of a student research team at Stanford Law School conducted as part of a Stanford Law School policy practicum taught in Fall 2025 by Jonathan J. Wroblewski and Robert Weisberg. We are grateful to the Law and Policy Lab at Stanford Law School, and the many people who helped in the creation of this report.

We also wish to extend our sincere gratitude to the Council on Criminal Justice’s (CCJ) Task Force on Artificial Intelligence (the Task Force), its members and advisors, and especially Jesse Rothman, Director of the Task Force, for their invaluable partnership and guidance throughout this project. The Task Force’s leadership in developing evidence-based standards and recommendations for the safe, ethical, and effective use of AI in the criminal justice system has been instrumental in shaping the lab’s analysis and recommendations. We deeply appreciate CCJ’s commitment to thoughtful governance of AI in criminal justice and for the opportunity to contribute to its ongoing efforts to build robust and actionable frameworks for AI use across criminal justice institutions.

Table of Contents

Executive Summary

Artificial intelligence (AI) is rapidly embedding itself into the core machinery of the criminal justice system, powering everyday decisions from police analysis of digital evidence and pattern detection in crime data to prosecutorial discovery management, charging recommendations, algorithmic risk assessments in courts, and large language models for summarizing records and drafting documents. These tools promise efficiency gains—processing vast data volumes, reducing backlogs, and optimizing scarce resources in an overburdened system—but they also carry profound risks: embedding bias, producing opaque or unreliable outputs, shifting unmonitored power to vendors, and influencing high-stakes liberty decisions like arrests, detention, sentencing, and release. The central problem is that AI capabilities are deploying without sufficient understanding of their mechanics, failure modes, or implications for constitutional rights, democratic accountability, and system legitimacy. Criminal-justice entities who encounter AI tools—thousands of under-resourced police departments, prosecutors' offices, courts, and probation units—lack the technical expertise to evaluate these tools rigorously, while vendors market directly to practitioners. This creates a governance gap: even well-intentioned actors cannot reliably apply emerging standards amid rapid technological changes, risking uneven, superficial oversight that undermines public trust.

Objectives of this White Paper

This White Paper seeks to bridge that gap by synthesizing institutional design principles for AI governance in criminal justice. It examines how other policy domains have addressed similar challenges of expertise asymmetry, evolving technologies, and high-stakes accountability. Our research formulates a twofold question: (1) What existing government entities could model—or directly assume—AI oversight roles? (2) How do those entities measure against desirable

attributes, revealing strengths, weaknesses, and hybrids? Starting from the premise that fragmented agencies need pooled support to properly govern the use of AI in criminal justice, we initially leaned towards a Federal governance model. We now have reconsidered this and believe that multi-level (Federal, State, local) governance with coordination may be the most feasible in today's politics. The goal is to equip policymakers at all levels with a framework to stand up durable institutions that operationalize the CCJ Task Force on AI's standards, ensuring AI enhances rather than erodes justice.

Institutional Design Criteria

We evaluate governance models across six core institutional design criteria that are consistent with CCJ Task Force's principles, balancing creation ease with sustained impact:

- **Startup Feasibility**: Legal/political/administrative ease of launch (e.g., executive action vs. statute), funding needs, and ramp-up time, prioritizing speed for urgent AI risks without sacrificing viability.
- **Relevant Expertise**: Dual composition with decision-makers with AI/criminal justice knowledge (legal, ethical, community) along with staff for technical assessment, empirical monitoring, and evaluation.
- **Transparency and Accessibility**: Open processes (public meetings, data releases), input mechanisms (hearings, comments), and usable outputs to enable scrutiny, error-correction, and stakeholder trust.
- **Organizational Stability**: Continuity amid politics (term structures, budgets, legal basis) to build institutional memory, iterate rules, and endure political changes.
- **Policy Influence**: Impact of outputs such as binding rules, required agency responses, grants, or reputational standards that drive uptake beyond advice.

- **Responsiveness**: Agility for fast-evolving AI, including expedited reviews, paired with slow-path monitoring, such as periodic data-driven updates.

Institutional Models Examined

1. FACA-Style Federal Advisory Committees

FACA committees, created under the 1972 Federal Advisory Committee Act, are renewable advisory bodies with a tenure of at least two years that provide balanced, expert advice to the Executive Branch. Examples include the National Commission on Forensic Sciences (NCFS) and the FBI's Criminal Justice Information Services (CJIS) Advisory Policy Board. These advisory committees generate non-binding recommendations.

2. Legislative Agencies

Legislative agencies support Congress with non-partisan analysis on complex issues. Examples are the Government Accountability Office (GAO) and the now-defunct Office of Technology Assessment (OTA, 1972–1995). GAO issues oversight reports and mandates agency responses, while the OTA produced intensive research on various technological developments. The role of legislative agencies is to inform legislation via expert reports, testimonies, and audits on emerging technologies like AI.

3. Sentencing Commissions

Sentencing commissions are independent expert bodies that develop system-wide standards for criminal justice decision-making through empirical analysis and iterative review. Examples include the U.S. Sentencing Commission and state commissions in North Carolina and Minnesota. These commissions combine pluralistic stakeholder representation with standing research staff, operate through transparent procedures, and regularly update guidance in response to observed outcomes. Although their authority is typically limited to sentencing, sentencing

commissions illustrate how nonbinding or semi-binding guidelines can achieve influence and legitimacy in technically complex domains.

4. **Congressional Committees**

Congressional committees encompass standing committees, subcommittees, and joint committees with direct congressional oversight. Examples include House/Senate Judiciary Committees and subcommittees, along with joint committees like the Joint Committee on Taxation (JCT). Their role involves conducting hearings, investigations, and bill drafting to shape policy on legal and constitutional matters.

Comparative Evaluation of the Models

No single institutional model performs well across all six design criteria. Instead, each reflects a different tradeoff among startup feasibility, expertise, stability, influence, and responsiveness. Advisory and legislative models are easiest to activate and can rapidly elevate AI governance on the policy agenda, but they often struggle to sustain deep technical capacity or long-term oversight. More durable expert bodies, by contrast, tend to accumulate legitimacy and analytic depth over time, but face higher political and administrative barriers to creation and may be slower to respond to rapidly evolving technologies.

FACA-style advisory committees and congressional subcommittees score highly on startup feasibility and transparency. They can be created quickly, convene diverse stakeholders, and generate public-facing guidance or hearings that shape early norms. However, their advisory nature, exposure to political turnover, and limited staffing constrain their ability to conduct sustained technical evaluation or monitoring. As the experience of the National Commission on Forensic Science illustrates, even highly effective advisory bodies may be short-lived, while congressional subcommittees remain vulnerable to shifting leadership priorities.

Legislative agencies and joint congressional committees offer deeper analytic capacity and greater institutional continuity, especially when supported by professional staff and formal reporting obligations. Entities such as the GAO demonstrate how nonbinding recommendations can nonetheless exert real influence through mandatory agency responses and reputational authority. At the same time, these models operate far upstream from day-to-day criminal justice practice and often trade speed for rigor, limiting their ability to provide timely, use-case-specific guidance in a fast-moving AI environment.

Sentencing commissions stand out as a hybrid model that balances expertise, stability, transparency, and iterative governance. Their experience shows how pluralistic decision-making bodies paired with standing research staff can translate technical analysis into durable standards that shape practice over time, even without broad enforcement authority. While their jurisdiction is narrower and their creation more demanding, sentencing commissions illustrate how AI governance institutions can earn legitimacy, adapt empirically, and persist across political cycles—suggesting that the most effective approach to AI governance in criminal justice may draw selectively from multiple models rather than replicating any single one wholesale.

The synthesis of these models reveals several consistent design principles. Durable AI governance requires permanence rather than ad hoc or temporary arrangements, "two-tier" expertise that pairs representative decision-makers with standing technical and research staff, and transparency that is structurally embedded rather than left to institutional culture. Influence need not depend on formal binding authority; mechanisms such as required agency responses, conditional funding, procurement standards, and credible technical evaluations can drive meaningful uptake even from advisory bodies. Responsiveness should combine rapid-response

mechanisms for emerging capabilities with slower, deeper processes for auditing and post-deployment monitoring.

These findings push toward a hybrid institutional approach that borrows selectively from multiple models rather than replicating any single one. Advisory committees and congressional subcommittees offer speed and agenda-setting power but struggle with durability and sustained technical depth. Legislative agencies provide analytic rigor and institutional continuity but operate too far upstream from day-to-day practice. Sentencing commissions come closest to balancing expertise, stability, transparency, and iterative governance, but their creation demands greater political investment. The most effective path forward will likely combine elements across these models—matched to the appropriate level of government—to close the growing gap between AI's pace of deployment and the justice system's capacity to govern it responsibly.

I. Introduction

Artificial intelligence is no longer a distant or speculative technology in the criminal justice system. It is becoming part of its everyday machinery. Police departments use algorithmic tools[1] to analyze digital evidence,[2] look for patterns in crime data,[3] and draft police reports.[4] Prosecutors rely on AI software to manage discovery[5] and support charging decisions.[6] Courts encounter algorithmic risk assessments[7] and large language models that promise to summarize records, draft documents, and assist with legal analysis.[8] Across the system, AI is being woven into decisions that affect liberty itself.

These technologies bring real promise. When they are carefully designed and responsibly governed, AI systems can help institutions use scarce resources more effectively, improve consistency, reduce backlogs, and expand access to legal and administrative services. They make

---

[1] Elizabeth E. Joh, Police Technology Experiments, 125 Colum. L. Rev. (Feb. 2025), https://www.columbialawreview.org/content/police-technology-experiments/.

[2] Susan Smalley, Law Enforcement Is Using AI to Synthesize Evidence. Is the Justice System Ready for It?, The Record (Sept. 29, 2025), https://therecord.media/law-enforcement-ai-platforms-synthesize-evidence-criminal-cases.

[3] Philip Lukens, Turning Data Into Decisions: Generative AI for Investigations and Intelligence, Police1 (Oct. 27, 2025), https://www.police1.com/vision/turning-data-into-decisions-generative-ai-for-investigations-and-intelligence.

[4] Fair & Just Prosecution, AI-Generated Police Reports: High-Tech, Low Accuracy, Big Risks (June 2025), https://fairandjustprosecution.org/wp-content/uploads/2025/06/AI-Generated-Police-Reports-High-Tech-Low-Accuracy-Big-Risks-June-2025.pdf.

[5] Digital Evidence Solutions for Public Defenders, NiCE, https://www.nicepublicsafety.com/justice/defense?_gl=1*778gig*_up*MQ..*_ga*MTg4NjEzNjgzNi4xNzY2MTAzMjk0*_ga_YYF02RP6R0*czE3NjYxMDMyOTMkbzEkZzAkdDE3NjYxMDMyOTMkajYwJGwwJGgxMjg4NjkxOTky.

[6] Mike Holley, An AI Primer for Prosecutors on Its Peril and Potential, Tex. Dist. & Cnty. Att'ys Ass'n: The Tex. Prosecutor (July–Aug. 2024), https://www.tdcaa.com/journal/an-ai-primer-for-prosecutors-on-its-peril-and-potential/.

[7] Matt Henry, Risk Assessment: Explained, The Appeal (May 25, 2019), https://theappeal.org/risk-assessment-explained/.

[8] Press Release, Grassley Releases Judges' Responses Owning Up to AI Use, Calls for Continued Oversight and Regulation (Oct. 23, 2025) U.S. Senate Committee on Judiciary, https://www.judiciary.senate.gov/press/rep/releases/grassley-releases-judges-responses-owning-up-to-ai-use-calls-for-continued-oversight-and-regulation.

it possible to process volumes of data and documentation that would be impossible to manage by hand. In a system that is chronically under-resourced and overburdened, those capabilities are understandably attractive.

At the same time, the risks are serious. Algorithmic systems can embed and amplify bias, operate in ways that are hard to see or understand, generate unreliable or misleading outputs, and shift decision-making power in directions that are difficult to monitor. Errors may not be obvious or isolated; they may be subtle, systematic, and persistent—especially for users without technical expertise. And the stakes are incredibly high, as these tools increasingly influence who is stopped, searched, charged, detained, sentenced, supervised, or released. When they fail, the consequences reach beyond efficiency and accuracy to constitutional rights, democratic accountability, and the legitimacy of the justice system itself.

As well-intentioned as policymakers and criminal justice practitioners including lawyers, judges, and police officers may be, without additional guidance, the lack of sufficient understanding of AI technology and its development may undermine their ability to effectively apply the Task Force's standards to specific use cases in a rapidly evolving AI landscape. For example, a judge may want to use a large language model (LLM) to help interpret a provision of a contract but may not fully understand the way the LLM works, how it was engineered, the risks of unreliable results, bias and other concerns. Police departments may want to use an AI product to process, analyze, or transcribe body camera video and audio but similarly may lack the expertise to identify the risks, reliability, or security around the product.[9]

[9] For instance, the body-camera company Axon has produced an AI tool called Draft One, which uses body camera footage to write police reports. See Draft One, Axon, https://www.axon.com/getstarted/draft-one?utm_source=ga&utm_medium=search&utm_campaign=d1_uv&utm_id=us_bb_q324&gad_source=1&gad_campaignid=21586001415&gbraid=0AAAAAD4QdgjxOc5MbDcoLz26yAfYV7sKM&gclid=CjwKCAiA3-3KBhBiEiwA2x7FdH36fCawPcbMB_adN-LE2GijBVeXl62fP-EQK4Kf42VJeZhjKHGs1BoCXlgQAvD_BwE (last visited Jan. 3, 2026). Dan Clark, a prosecutor in Seattle's King County Prosecutor's Office, spoke with us about

Artificial intelligence capabilities and products are being developed and deployed in criminal justice without sufficient understanding of how they work, attention to their failure risks, or analysis of how they implicate constitutional and other values.

The central problem this White Paper addresses can therefore be stated as follows:

Artificial intelligence capabilities and products are being developed and deployed in criminal justice without sufficient understanding of how they work, attention to their failure risks, or analysis of how they implicate constitutional and other values. Judges, prosecutors, police officers, probation officers, and policymakers are often asked to evaluate and adopt complex technologies without the technical background needed to assess their design, limitations, or possible failures. While some criminal justice agencies will have the capacity to put together multi-disciplinary teams to evaluate AI technology, most will not. At the same time, vendors market these systems directly to practitioners, who may lack both the time and the expertise to conduct meaningful independent evaluation. The result is a growing gap between the pace of technological change and the capacity of institutions to govern it responsibly.

Even defining the object of governance is not straightforward. "Artificial intelligence" is not a single technology but a shifting collection of methods—including machine learning, pattern recognition, natural language processing, computer vision, and automated decision-support systems. Some of these tools are genuinely novel and opaque; others resemble familiar statistical

---

his decision not to accept AI-generated police reports, which frustrated many of the 39 police agencies in King County. While acknowledging police officers' legitimate and well-meaning aim of reducing time spent writing reports so as to increase time policing, Clark believed concerns about the tool outweigh its benefits. He cited concerns about the product's data collection—how it handles translation, tone, accents, multiple speakers, radio noise in the background—as well as its accuracy—Axon refused to share hallucination rate data with Clark—and the design of the product—Draft One "neutralizes" biased language, raising concerns that the reports it generates are inaccurate, and may cover up police bias and misconduct. Clark was further concerned that officers may behave differently in reliance upon the AI tool. In particular, he worried officers may neglect to report things they observe through their other senses, like smells, or that they may start speaking differently so their body cameras and the AI tool would be able to capture more information, an especially problematic practice in cases involving violence, where, for example, remarking aloud on a victim's injuries may be wildly inappropriate.

techniques. In practice, it is often unclear when the use of an algorithm becomes the use of "AI" for governance purposes. Does software that flags anomalous financial transactions in a white-collar investigation count as AI, or is it just advanced statistical analysis? Does an algorithm that redacts information to enable race-blind charging fall within the scope of AI governance even if no branded "AI product" is involved? These ambiguities make regulation harder and suggest that governance cannot rest on technical labels alone.

For that reason, this White Paper focuses on general principles for governing AI capabilities at large, rather than particular AI products. A purely product-centered approach is likely to be both too narrow and quickly outdated, as specific tools evolve, are replaced, or are simply rebranded. A capability-centered approach instead asks what a system is doing—automated pattern detection, risk scoring, language generation, biometric identification, predictive classification—and what kinds of decisions or processes those functions affect. This framing makes it possible to track impact and function rather than vendor-specific implementations, and to remain relevant as technologies change. Developed in partnership with the Council on Criminal Justice's Task Force on Artificial Intelligence, this White Paper builds upon the Task Force's efforts[10] to establish governing principles, procurement standards, and case studies for safe, ethical, and effective use of AI in criminal justice. It extends that work by examining potential structures and mechanisms through which Congress, the Executive Branch, State and local governments, and non-governmental organizations too, could institutionalize the Task Force's recommendations, and provides a foundation for assessing durable governance frameworks that can translate expert recommendations into actionable policies and implement them.

[10] Council on Criminal Justice, Task Force on Artificial Intelligence, https://councilonj.org/artificial-intelligence/; Principles for the Use of AI in Criminal Justice (Oct. 2025) https://councilonj.org/principles-for-the-use-of-ai-in-criminal-justice/.

To synthesize recommendations for governing the use of artificial intelligence in criminal justice, this White Paper looks at how other complex and evolving policy domains have dealt with similar problems of expertise, legitimacy, flexibility, and accountability. By examining existing institutional models, it aims to identify institutional design principles that could support AI governance in criminal justice. **Our goal is not to prescribe a single institutional blueprint, but to clarify the tradeoffs and give policymakers a framework for thinking through their choices to stand-up an AI governance entity to apply the Task Force standards to individual AI use cases**.

Although this White Paper does not advocate a single institutional blueprint, it starts from the practical reality that many criminal justice entities will not be well positioned to apply even a well-designed evaluation framework on their own. The criminal justice system is highly fragmented, comprising thousands of police departments, prosecutor's offices, courts, probation agencies, and administrative bodies, many of which are small, under-resourced, and lacking in-house technical or analytic capacity. Even where leadership is attentive to the risks posed by artificial intelligence, evaluating the design, training data, performance characteristics, and failure modes of AI products requires expertise and time that most agencies simply do not have. Vendors, by contrast, are often sophisticated, well-resourced, and highly effective at marketing complex tools in ways that obscure limitations or shift attention away from downstream risks. As a result, placing primary responsibility for AI governance on individual agencies risks producing uneven, duplicative, and often superficial review—undermining the Task Force's goal of responsible, consistent, and values-aligned deployment across the system. In practice, most criminal justice entities lack not only the technical capacity to rigorously evaluate AI products, but also the institutional distance needed to assess them dispassionately. Because many AI tools promise

immediate, localized benefits—such as reduced workloads, faster processing, or budgetary savings—agencies may discount longer-term risks, externalities, or constitutional concerns that fall outside their immediate operational interests.

This capacity gap suggests a potential role for specialized expert entities that can aggregate expertise, conduct sustained analysis, and provide authoritative guidance that individual agencies can rely on when evaluating AI use cases. Such entities need not be limited to novel government agencies. In addition to Federal, State, or local public bodies, there may be important roles for non-governmental or quasi-governmental institutions affiliated with existing professional or intergovernmental associations, such as organizations representing prosecutors, attorneys general, governors, or court administrators. These bodies may be able to convene technical expertise, develop shared evaluation resources, issue best-practice guidance, or facilitate coordination across jurisdictions in ways that are more feasible than agency-by-agency assessment. Recognizing this broader institutional landscape helps clarify why some form of centralized or pooled expertise—whether governmental, non-governmental, or hybrid—may be an important complement to a framework that is ultimately meant to guide decision-making at the individual agency level.

Taking as a starting point that some government entity is needed to provide guidance to criminal justice policymakers and practitioners, we formulated a twofold research question:

1) What government entities currently exist that could either serve as models for a novel agency dedicated to AI in criminal justice, or could take on this work themselves?
2) What attributes are desirable for an AI in criminal justice government entity, and, using a set of criteria we have developed, to what extent do each of the models identified in (1) exhibit these attributes, and how do the models compare to one another?

Motivating this approach to our research is the conviction that it would be ideal if an already existing government agency could provide guidance on artificial intelligence in criminal justice, but failing that, we sought to identify various models for such an agency and develop a framework to compare their respective strengths and weaknesses. This would be ideal because existing agencies already possess statutory authority, administrative infrastructure, and institutional legitimacy, allowing guidance to be developed and disseminated more quickly and with greater likelihood of uptake. Leveraging an established entity also reduces startup costs and political barriers, while minimizing duplication in a policy area where time, expertise, and credibility are especially scarce.

At the outset of the project, we assumed that a single entity at the Federal level would be preferable over every state having its own artificial intelligence and criminal justice oversight body. Our research led us to rethink that initial assumption, particularly as state and/or local agencies may be far more feasible to establish in the current political climate than a novel Federal entity. We now believe that the institutional templates we have researched could be shared as part of the Task Force's final work product and could be used to stand up entities at the Federal, State, and/or local levels to evaluate and advise stakeholders on granular AI use cases in the criminal justice system on an ongoing basis. We recognize that certain models may be more or less applicable to state-level policy making, and if many of these governance entities in fact are developed, we think that some multi-level coordination will be helpful for efficiency, consistency, and identification and utilization of best practices.

The remainder of the White Paper is organized as follows. Part II discusses the institutional design criteria that guide the analysis of institutional models in the coming section. Part III surveys the institutional models that we considered as potential templates for governing AI in criminal

justice, including FACA-style Federal advisory committees (such as the National Commission on Forensic Science), the U.S. Sentencing Commission, the North Carolina Sentencing and Policy Advisory Commission, the Minnesota Sentencing Guidelines Commission, joint congressional committees, Judiciary Committee subcommittees, the former Office of Technology Assessment, and the Government Accountability Office. Part IV synthesizes the lessons from these models along three dimensions: (a) core design features such as independence, expertise, transparency, responsiveness, and stability; (b) mechanisms for policy impact, including binding and nonbinding recommendations, conditional grants, and technical reporting; and (c) the relative strengths and weaknesses of Federal, State, and multi-level governance. Part V concludes by summarizing the implications of this analysis and identifying key considerations for future efforts to institutionalize AI governance in criminal justice.

II. Institutional Design Criteria for AI Governance Bodies

This White Paper looks at institutional models based on six core criteria for the design of AI governance entities: startup feasibility, relevant expertise, transparency and accessibility, organizational stability, policy influence, and responsiveness. Each criterion captures a distinct dimension of institutional performance and taken together they provide a structured basis for evaluating and comparing alternative institutional models.

A. Startup Feasibility

Startup feasibility refers to the legal, political, and administrative ease with which an institution can be created, staffed, and operated. It includes the formal preconditions for

establishment (e.g., statutory enactment versus internal reorganization), the dependence on funds, and the time required to reach functional capacity.

The purpose of foregrounding startup feasibility is to balance institutional creation and governance abilities. Arrangements such as subcommittees, advisory boards, or task forces that are within existing institutional bodies can be constituted rapidly, and are well-suited to agenda-setting, fact-finding, and early norm articulation in a fast-moving AI environment. In contrast, more robust entities created by statute and backed by durable funding streams are harder to establish but are better positioned to sustain iterative governance functions over years or decades. Treating feasibility as a design variable rather than an afterthought clarifies when policymakers should prioritize speed and reversibility and when they should instead invest in institutional forms that are slower to establish but more resistant to short-term political shocks.

B. Relevant Expertise

Relevant expertise captures both the composition of the decision-making body and the nature of its standing staff. It includes: (i) whether members possess substantive knowledge of the regulated domain and affected systems (which in this case would be expertise in AI and in criminal justice, and adjacent fields) and (ii) whether the institution maintains analytic capacity for technical assessment, empirical research, evaluation, and monitoring.

The point of isolating this factor is to emphasize that "expertise" is not a scalar attribute but a two-tier requirement. For an AI governance institution to be credible and legitimate, it must combine the integration of legal, ethical, and community perspectives with rigorous technical work. Design choices about appointment mechanisms, staffing models, and resourcing thus directly determine whether the institution can (a) interpret complex AI systems, (b) evaluate

impact over time, and (c) reflect the values and experiences of those most affected by AI in the criminal justice system.

C. Transparency and Accessibility

Transparency and accessibility refer to the degree to which an institution's processes, deliberations, outputs, and underlying data are open to observation, scrutiny, and participation by external stakeholders. This encompasses formal obligations of openness (such as public meetings, published reports, and open data releases), opportunities for structured input (hearings, notice-and-comment mechanisms, participatory procedures), and practical accessibility (the intelligibility and usability of outputs to relevant stakeholders).

Incorporating transparency as an explicit design criterion reflects the reality that in AI governance, openness is not merely a normative virtue but a functional precondition for trust, error-correction, and iterative learning. AI systems are often opaque, and many failures or disparate impacts are only observable ex post. Institutions that structurally embed transparency and stakeholder access are better placed to mobilize external expertise, enable independent research, and create feedback loops between regulators, implementers, and affected communities.

D. Organizational Stability

Organizational stability captures the extent to which an institution can maintain continuity of operations, mission, and expertise over time, notwithstanding changes in political leadership, funding priorities, or external pressures. Key determinants include the legal basis of the entity, the structure and staggering of terms, budget arrangements, and the presence of routine procedures that anchor ongoing work.

Highlighting stability as a separate design factor underscores that durable AI governance depends on more than episodic interventions. Effective oversight of complex, evolving

sociotechnical systems requires institutional memory, empirical research, and repeated cycles of rulemaking, evaluation, and revision. Bodies that can be readily dissolved, reorganized, or starved of resources will struggle to sustain these functions, regardless of the quality of their initial mandate. Treating stability as a design constraint pushes designers toward arrangements that can accumulate expertise and credibility across electoral cycles and shifting political coalitions.

E. Policy Influence

Policy influence refers to the mechanisms through which an institution affects substantive outcomes. This includes formal authority (e.g., binding rules, guidelines with legal force), structured uptake pathways (such as statutory requirements that agencies respond to recommendations or explain deviations), and reputational capital that can lead other actors to treat outputs as de facto governance standards.

The reason to foreground policy influence is that an institution's impact is not determined solely by whether its outputs are formally binding. Advisory and analytic bodies can materially shape practice when their recommendations are systematically integrated into procurement rules, funding conditions, internal agency policies, or professional norms. Conversely, even entities with nominal rulemaking power may have limited real-world effect if their outputs are ignored in implementation or circumvented in practice. Incorporating this criterion forces attention to the design of concrete uptake mechanisms, such as mandatory response requirements, reporting obligations, and linkages to enforcement or budgetary levers that convert guidance into changes in organizational behavior.

F. Responsiveness

Responsiveness denotes an institution's capacity to detect, prioritize, and address emerging developments, risks, and failures in a timely manner. It includes both temporal dimensions (speed

of reaction, frequency of review) and substantive dimensions (ability to redirect research, update standards, and recalibrate interventions as new developments take place).

Responsiveness as a design dimension reflects the dual temporal demands of AI governance. Rapid development of novel capabilities and harms may require expedited hearings, guidance, or interim controls, whereas the effect on aspects such as sentencing patterns, cumulative bias, or institutional dependence on automated tools emerge slowly and can only be addressed through sustained monitoring and periodic revision. A well-designed AI governance architecture therefore needs complementary fast and slow pathways where mechanisms for rapid response are coupled with deeper, data-driven assessment and rule-updating processes. Embedding responsiveness institutionally through regular amendment cycles, scheduled reviews, and ongoing data collection reduces reliance on ad hoc political attention.

Taken together, these six criteria function as a structured framework for analyzing and comparing institutional design options for AI governance in the criminal justice system. They highlight the trade-offs that are otherwise easy to obscure, such as speed of formation versus durability, breadth of representation versus depth of technical capacity, openness versus vulnerability, formal authority versus practical uptake, and rapid reaction versus considered, iterative adjustment.

III. Institutional Models

A. FACA Model

This chapter considers whether a Federal advisory committee could effectively ensure that AI deployment in criminal justice is consistent with fundamental Constitutional principles and the broader principles developed by the Task Force. After evaluating the general effectiveness of the

Federal Advisory Committee Act (FACA) model, it considers two case studies: the Department of Justice's National Commission on Forensic Science (NCFS) and the FBI's Criminal Justice Information Services (CJIS) Advisory Policy Board.

The analysis suggests that the FACA model might indeed be a promising option for AI governance, with the caveat that the success of the model appears context-dependent. As our case studies illustrate, differing committee structures and emphases can lead to differing results—the NCFS' success was broad, but limited by political instability, and the CJIS' continued influence is confined to a narrow sector.

The Federal Advisory Committee Act (FACA)

In 1972, Congress passed FACA to formalize the creation and maintenance of Federal advisory committees, which had historically advised the Executive Branch.[11] FACA's statutory text directed newly created advisory committees to prioritize "efficiency, balance, and openness."[12] However, as described below, the FACA-model has inherent limitations; its emphasis on efficiency can limit the long-term stability of a governing body, and prioritizing balance may limit agility.

A. Startup Feasibility

One of the greatest strengths of a FACA-style committee is its start-up feasibility.[13] Multiple actors have the ability to create a committee—Congress, the President, and agency heads

[11] 5 U.S.C. Appendix; *see also* WENDY GINSBERG & CASEY BURGAT, CONG. RSCH. SERV., R44232, CREATING A FEDERAL ADVISORY COMMITTEE IN THE EXECUTIVE BRANCH 1 (2016); Michael J. Mongan, *Fixing FACA: The Case for Exempting Presidential Advisory Committees from Judicial Review Under the Federal Advisory Committee Act,* 58 STAN. L. REV. 895, 896-903 (2005).

[12] Mongan, *supra* note 1, at 903.

[13] These relative low-barriers to committee creation may help explain the proliferation of Federal advisory committees; in the status quo, there are nearly 1,000 in operation. *See* U.S. GEN. SERVS. ADMIN., FEDERAL ADVISORY COMMITTEE ACT (FACA) DATABASE (2025), https://www.facadatabase.gov/FACA/s/.

all have authority to independently create a committee.[14] There is also broad latitude as to what counts as a committee; the statutory text broadly defines a Federal advisory body as any "committee, board, commission…task force…or any subcommittee" that provides some advice to the Executive Branch and includes at least one non-Federal government employee.[15] Once created, new advisory committees must only meet a few requirements: they must publish a charter which establishes the committee's mission and bylaws, and must apply for Congressional re-chartering every two-years.[16] Together, these limited requirements mean that the start-up costs for Federal advisory committees is relatively low; action by a single political body, like Congress or the Department of Justice, would be enough to create an AI-focused advisory committee.

B. Relevant Expertise

The FACA model also inherently prioritizes inclusion of diverse perspectives. As noted, FACA's statutory text puts an emphasis on "balance."[17] In practice, this means Federal committee membership can prioritize multiple types of diversity, including regional diversity, diversity across different work sectors, or representation from multiple governmental branches and agencies.[18] In the AI policy context, this might mean recruiting stakeholders from across the criminal justice *and* AI sectors. Importantly, beyond an emphasis on balance, the typical FACA committee structure might enable participation by non-traditional government actors. Many Federal advisory committees meet infrequently, some only quarterly or bi-annually.[19] These relatively infrequent

---

[14] GINSBERG & BURGAT, *supra* note 1, at 1.
[15] 5 U.S.C. §1001(2)(A).
[16] CONG. RSCH. SERV., R4453, FEDERAL ADVISORY COMMITTEES: AN OVERVIEW AND INTRODUCTION 4-9 (2016)
[17] 5 U.S.C. Appendix §5(2).
[18] CONG. RSCH. SERV., R47894, The FEDERAL ADVISORY COMMITTEE ACT: OVERVIEW AND CONSIDERATIONS FOR CONGRESS 15 (2024).
[19] *See, e.g.*, U.S. FOOD & DRUG ADMIN., QUESTIONS AND ANSWERS REGARDING ADVISORY COMMITTEE MEMBERSHIP (2018), fda.gov/advisory-committees/advisory-committee-membership/questions-and-answers-regarding-advisory-committee-membership (claiming committees typically meet less than four times annually); U.S. GEN. SERVS. ADMIN., FACADATABASE.GOV: NATIONAL INSTITUTE OF CORRECTIONS ADVISORY BOARD MEETINGS

meetings enable professional industry stakeholders, like those in the technology industry, to serve as committee members while working elsewhere full-time.[20] The FACA model thus allows for diverse stakeholders to convene. This emphasis on balanced membership closely tracks the Task Force's concern that AI governance should not be dominated by either vendors or end-users alone, but instead reflect the perspectives of technologists, criminal justice practitioners, legal experts, and impacted communities. In this respect, the FACA statute's balance requirement maps unusually well onto the Task Force's normative commitments.

C. Transparency and Accessibility

The statutory language of FACA also ensures any Federal advisory committees are largely transparent and publicly accessible. By law, Federal advisory committees must make committee meetings publicly available and provide notice of such meetings.[21] Committees must also make minutes, transcripts, and other committee-generated documents publicly accessible.[22] Together, these requirements ensure strong public oversight of all Federal advisory committees. There are also pre-established Congressional oversight mechanisms; all FACA committees must be re-charted by Congress every two years, enabling regular Congressional oversight.

D. Organizational Stability

The long-term stability of FACA-style committees can be limited or subject to changing political environments. As noted above, by statute, Federal advisory committees have a two-year

(2025), https://www.facadatabase.gov/FACA/s/meeting-members-advisory-reports?recordId=a10SJ000004FZETYA4&tabset-29055=1 (noting only three annual meetings for a DOJ-housed advisory committee).

[20] *See, e.g.*, U.S. GEN. SERVS. ADMIN., FACADATABASE.GOV: INDUSTRY TRADE ADVISORY COMMITTEE ON DIGITAL ECONOMY MEMBERS (2025), https://www.facadatabase.gov/FACA/s/meeting-members-advisory-reports?recordId=a10t0000001gzkOAAQ (noting diverse committee membership, including representatives from Apple, Google, and HP).

[21] Mongan, *supra* note 1, at 904.

[22] *Id.*

tenure; if they are not re-charted at the end of that time, they cease to exist.[23] As a result, as our case studies demonstrate, the long-term stability and viability of a committee can vary widely. In some cases, a committee may never be extended, capping its influence at two years.

E. Policy Influence

The policy influence of a singular commission can range widely, as noted in our case studies below, and may interact directly with the longevity or focus of a singular committee. And importantly, under FACA, any recommendations made by advisory committees are non-binding and strictly *advisory*.[24] With that said, there is high potential for Federal advisory committees to be impactful on agencies; as noted, both the NCFS and CJIS had most of their recommendations adopted and implemented.[25] However, other committees have been significantly less effective. Survey data by the General Services Administration (GSA) from FY 2025 note that only 38% of committees self-reported contributing to "major policy changes."[26] Some academics have been similarly critical of the influence of committees, arguing they are primarily utilized as an ideological "counterweight" against the views of longstanding agency employees.[27]

F. Responsiveness

Like other variables, responsiveness is committee specific, dependent on the meeting frequency and productivity of a committee's members. As noted,[28] many Federal advisory committees meet only a few times a year. Accordingly, committees that meet infrequently may

---

[23] *Id.*

[24] GINSBERG & BURGAT, *supra* note 1, at 1.

[25] *See* U.S. GEN. SERVS. ADMIN., FACADATABASE.GOV: 2025 CURRENT FISCAL YEAR REPORT: CRIMINAL JUSTICE INFORMATION SERVICES ADVISORY POLICY BOARD (2025), https://www.facadatabase.gov/FACA/apex/FACACommitteeLevelReportAsPDF?id=a10t0000001gzqJAAQ (noting 96% recommendation adoption rate by DOJ)*; see also supra* notes 30-31 and accompanying text.

[26] U.S. GEN. SERVS. ADMIN., FACADATABASE.GOV, FY 2025 PERFORMANCE MEASURES (ACR) TOTALS (2025), https://www.facadatabase.gov/FACA/s/PerformanceMeasureTotals.

[27] Brian D. Feinstein & Daniel J. Hemel, *Outside Advisers Inside Agencies*, 108 GEO. L. J. 1139, 1145 (2020).

[28] *See supra* note 9.

not be able to quickly generate recommendations, something that may prove especially problematic in a fast-moving area like AI. With that said, many committees, despite infrequent meetings, still produce significant recommendations. The GSA reports that in FY 2025, currently chartered Federal advisory committees had generated over four million recommendations over the course of their committees' lifetimes.[29] Thus, while infrequent meetings don't necessarily limit committees' ability to *produce* work, they may limit their ability to act quickly.

Case Studies

To further evaluate the potential effectiveness of the FACA model, we examined two different Federal advisory committees, the National Commission on Forensic Science (NCFS) and the Criminal Justice Information Service (CJIS) Advisory Policy Board, highlighting specific variables which are most relevant to the Task Force's concerns. As described, these two committees had different structures and priorities, and ultimately, had differing results.

1. National Commission on Forensic Science (NCFS)

The NCFS was a joint advisory body established under FACA by the Department of Justice (DOJ) and the National Institute of Standards and Technology (NIST), operating from 2013 to 2017. Composed of 30 appointed members and 10 *ex officio* members, the Commission brought together experts from law enforcement, academia, the judiciary, and the forensic science community.[30] Its mandate was to advise and provide recommendations to the DOJ on the methods and strategies related to forensic science.[31]

---

[29] U.S. GEN. SERVS. ADMIN., FACADATABASE.GOV, FISCAL YEAR 2025 GOVERNMENT TOTALS (2025), https://www.facadatabase.gov/FACA/s/GovtWideTotals.

[30] U.S. DEP'T OF JUST. & NAT'L INST. OF STANDARDS & TECH., BYLAWS OF THE NATIONAL COMMISSION ON FORENSIC SCIENCE, 2 (2016).

[31] U.S. DEP'T OF JUST., CHARTER OF THE NATIONAL COMMISSION ON FORENSIC SCIENCE, 1 (2015).

A. Relevant Expertise

One of the Commission's most notable strengths was its accumulation of diverse expertise across criminal justice and forensic science. Members were deliberately selected by the Attorney General (AG), in consultation with NIST leadership, to ensure balanced representation from science, law, law enforcement, and academia, with *ex officio* participants further enriching deliberations.[32] Annual consultations preserved this balance over time,[33] while outreach through the Federal Register and professional societies ensured continued access to specialized expertise.[34] The NCFS further strengthened its expertise through the use of subcommittees established with the AG's approval.[35] These subcommittees focused on specialized areas within forensic science and criminal justice, allowing for deeper technical analysis than would have been possible at the full Commission level. This multidisciplinary and structured approach is particularly instructive for AI governance, where effective oversight requires the integration of technical, legal, and institutional perspectives.

B. Transparency and Accessibility

Transparency and accessibility were central to the NCFS's operations. In line with FACA requirements, Commission meetings were generally open to the public, with materials presented or discussed made available at the time of the meeting to ensure real-time access to deliberations.[36] Members of the public could attend meetings and, at the discretion of the Co-Chairs or Vice-Chairs, provide oral comments during designated public comment periods, while written

[32] *Id.* at 2.
[33] U.S. DEP'T OF JUST. & NAT'L INST. OF STANDARDS & TECH., *supra* note 3, at 2.
[34] U.S. DEP'T OF JUST., *supra* note 4, at 2.
[35] *Id.*
[36] U.S. DEP'T OF JUST. & NAT'L INST. OF STANDARDS & TECH., *supra* note 3, at 3.

submissions could be submitted to the Designated Federal Officer (DFO) at any time.[37] Meetings could be closed only where the DFO determined that public disclosure would be harmful to governmental, industry, or other interests, and only with advance approval from the AG and the DOJ's Office of General Counsel.[38] Overall, this framework ensured a high level of transparency and accessibility, reinforcing the legitimacy of the Commission's work and offering a strong procedural model for any commission designed to govern AI systems.

C. Organizational Stability

Despite these strengths, the NCFS faced significant limitations in organizational stability. Its charter included a two-year termination date, renewable under FACA, and the Commission ultimately operated for only four years before being phased out following a change in presidential administration.[39] The charter expired in April 2017, illustrating the vulnerability of FACA-based advisory bodies to political transitions. This limited lifespan constrained the Commission's ability to embed long-term reforms and raises concerns about whether a similar structure could provide sustained oversight for AI systems that evolve continuously and require ongoing governance.

D. Policy Influence

The NCFS's influence depended heavily on DOJ adoption of its advisory outputs and was further limited by its non-binding authority and short duration.[40] Its work products consisted of "Recommendations to the Attorney General," which provided formal advisory guidance, and "Views of the Commission," which reflected the group's consensus on policy or technical

---

[37] *Id.*
[38] *Id.* at 3-4.
[39] U.S. DEP'T OF JUST., *supra* note 4, at 2.
[40] Jules Epstein, The National Commission on Forensic Science: Impactful or Ineffectual, 48 SETON HALL L. REV. 743 (2018).

matters.[41] The NCFS made several credible recommendations, many of which DOJ adopted. The DOJ endorsed proposals on universal accreditation, laboratory best practices, improved access to fingerprint databases,[42] and directed labs to stop using the term "reasonable scientific certainty" or "'reasonable [discipline] certainty'' in reports and testimony.[43] However, the advisory nature of its authority and its finite duration limited the breadth and durability of its impact, highlighting a challenge for any future governance framework modeled on the NCFS.

Interviews with former NCFS participants underscore that the Commission's influence depended less on its formal advisory status than on a set of design choices that fostered internal legitimacy and external trust. Interviewees consistently emphasized three features in particular: (1) a clearly defined and normatively grounded charter that focused the Commission on concrete, system-wide problems rather than abstract principles; (2) the use of subcommittees to allow sustained, technically rigorous work outside plenary meetings; and (3) close, but not subordinate, institutional proximity to DOJ leadership, which ensured recommendations were salient without being pre-cleared. At the same time, interviewees noted that the Commission's vulnerability to political transition—and its lack of any mechanism tying recommendations to downstream adoption—ultimately limited its durability. These insights suggest that the success of a FACA-style AI governance body would turn less on the advisory label itself and more on whether its charter, internal structure, and institutional positioning were deliberately engineered to translate expertise into durable influence.

---

[41] U.S. DEP'T OF JUST. & NAT'L INST. OF STANDARDS & TECH. (n3), 3.

[42] LORETTA E. LYNCH, ATT'Y GEN. OF THE U.S., MEM. FOR HEADS OF DEPARTMENT COMPONENTS: RECOMMENDATIONS OF THE NATIONAL COMMISSION ON FORENSIC SCIENCE; ANNOUNCEMENT FOR NCFS MEETING NINE, 1 (Mar. 2016), https://www.justice.gov/ncfs/file/841861/dl?inline.

[43] LORETTA E. LYNCH, ATT'Y GEN. OF THE U.S., MEM. FOR HEADS OF DEPARTMENT COMPONENTS: RECOMMENDATIONS OF THE NATIONAL COMMISSION ON FORENSIC SCIENCE; ANNOUNCEMENT FOR NCFS MEETING ELEVEN, 1 (Sept. 2016), https://www.justice.gov/archives/opa/file/891366/dl?inline.

2. Criminal Justice Information Services (CJIS)

In 1994, the FBI created the Criminal Justice Information Services (CJIS) Advisory Policy Board (APB) to centralize digital information shared across various law enforcement agencies, like the National Crime Information Center or National Instant Criminal Background Check System.[44] Following its inception, the APB was statutorily codified and has remained in existence since, with consistent re-chartering.[45]

A. Relevant Expertise

One critical difference from the NCFS is that the APB has a unique emphasis on geographic, and not professional, diversity. While the Board has a total of 35 members, 20 of these members must be chosen from state and local law enforcement agencies from five distinct geographic regions (Western, North Central, Southern, Northeastern, and Federal).[46] Other Board members include industry professionals selected by the current FBI Director or who represent industry professional groups, like the National Sheriff's Association.[47] Notably, almost all of the APB is composed of individuals who work in law enforcement—it is not a sector-diverse group. This structure makes less sense for a policy challenge like AI governance, where input from diverse stakeholders, like those in technology *and* in the criminal justice sectors is needed.

B. Transparency and Accessibility

Like other FACA committees, the vast majority of the APB's work is publicly accessible. The APB typically meets twice a year,[48] and members of the public and law enforcement agencies

---

[44] U.S. FED. BUREAU OF INTEL., MANAGEMENT OF CJIS SYSTEM SHARED WITH CRIMINAL JUSTICE COMMUNITY (2014), https://le.fbi.gov/cjis-division/cjis-link/management-of-cjis-system-shared-with-criminal-justice-community.
[45] *See id.*; *see also* 28 C.F.R. §20.35.
[46] U.S. FED. BUREAU OF INTEL., ADVISORY POLICY BOARD INVOLVES USERS TO HELP MANAGE CRIMINAL JUSTICE SYSTEM (2022), https://le.fbi.gov/cjis-division/cjis-link/advisory-policy-board-involves-users-to-help-manage-criminal-justice-systems.
[47] *Id.*
[48] U.S. GEN. SERVS. ADMIN., 2025 CURRENT FISCAL YEAR REPORT, *supra* note 16.

are encouraged to submit ideas and proposals to the Board.[49] Like the NCFS, the APB demonstrates that advisory committees largely uphold their commitment to transparency.

C. Policy Influence

The APB has been uniquely successful at having its recommendations adopted; in FY 2025, the Board reported that 96% of its recommendations had been adopted by the FBI, and unlike the NCFS, the APB has remained influential for decades.[50] While the reasons for this success are not entirely clear, the APB serves a unique regulatory function, which incentivizes greater agency compliance. Unlike the NCFS, the APB has a unique "carrot" which encourages law enforcement (and the FBI) to follow its recommendations. The CJIS sets national standards for access to critical digital technologies based on APB recommendations. When law enforcement agencies fail to comply with these standards, they lose access to the technology.[51] This leverage gives the CJIS the unique ability to influence its target audience, the greater law enforcement community. An AI-focused advisory committee could approximate this form of leverage by tying its guidance to access-dependent systems or benefits rather than relying solely on persuasion. For example, a committee could be structured so that compliance with its standards becomes a condition for eligibility for Federal grants, access to shared AI evaluation tools, inclusion on approved procurement lists, or participation in interoperable data systems. Even without formal rulemaking authority, these linkages could create meaningful incentives for adoption while preserving the

[49] U.S.FED. BUREAU OF INTEL., THE CJIS ADVISORY PROCESS (2025), https://le.fbi.gov/cjis-division/the-cjis-advisory-process.

[50] U.S. GEN. SERVS. ADMIN., 2025 CURRENT FISCAL YEAR REPORT, *supra* note 16; *see also* U.S. FED. BUREAU OF INTEL., MANAGEMENT OF CJIS SYSTEM, *supra* note 36.

[51] *Compare* U.S. DEP'T OF JUST., CRIMINAL JUSTICE INFORMATION SERVICES: SECURITY POLICY (2020), https://arjis.org/RegionalPolicies/CJIS_Security_Policy_v5-9_20200601.pdf, *with* CAL. DEP'T OF JUST., FBI CRIMINAL JUSTICE INFORMATION SERVICES (CJIS) SECURITY POLICY, VERSION 5.9.5 (2024), https://oag.ca.gov/system/files/media/2024-isrs-002.pdf (suggesting California law enforcement will lose access to CJIS' services if they fail to comply with updated CJIS policy).

advisory character of the committee. The CJIS example thus illustrates that policy influence under FACA need not be binary; carefully designed institutional linkages can substantially increase uptake without converting an advisory body into a regulator. While it may be possible to envision implementing similar access standards in an AI-technology space, it is worth noting the APB's high-level of success might be primarily due to this unique attribute.

Conclusion

Ultimately, the Federal Advisory Committee Act-style of Federal advisory commissions provides a potentially compelling framework for governing the use of AI in the criminal justice system. As an examination of the NCFS and CJIS demonstrates, many FACA-style commissions are able to convene diverse stakeholders, in a publicly accessible way, and generate potentially highly influential recommendations to Federal agencies efficiently. However, these case studies also illustrate that these positive characteristics are potential pitfalls. As the NCFS demonstrates, unpredictably changing political environments can result in the shuttering of a successful commission. Even when committees do achieve political longevity, like the CJIS, it is difficult to assess whether success is merely attributable to a committee-specific focus or structure. Thoughtful structuring—particularly around charter specificity, subcommittee design, and linkage to downstream incentives—could allow an AI-focused advisory committee to retain FACA's startup advantages while mitigating its traditional weaknesses in durability and policy influence.

B. Legislative Agency Model

Legislative agencies offer another framework for stewarding the Task Force's recommendations for AI in criminal justice. As with their peers in the executive branch, legislative agencies come in a variety of forms and serve distinct functions. They are unified by their accountability to Congress and their unique proximity to the legislative process. As the power of

the Federal government has expanded over the last century, legislative agencies have offered critical support to members of Congress legislating on problems of significant complexity.[52] The Congressional Budget Office, for example, retains a staff of budget experts and economists to help legislators understand the budgetary consequences of omnibus bills touching every corner of the American economy. AI governance—with its technical complexity and pressing policy implications—is the kind of issue where Congress may benefit from the sustained and rigorous focus of a dedicated legislative agency.

Case Studies

The menu of legislative agencies is broad, and no two models share precisely the same characteristics. For this White Paper, we have evaluated two entities: the Office of Technology Assessment (OTA) and the Government Accountability Office (GAO). As the OTA has been inactive since the mid-1990s, our analysis considers the possibility of reviving it or a similar institution focused on stewarding the Task Force's recommendations. The GAO, by contrast, continues to operate, and this section considers the desirability of folding the Task Force's recommendations into the agency's ongoing work.

Additionally, recent developments in administrative law suggest an impending rebalancing of power between the branches. With the demise of *Chevron* deference and the Trump Administration's challenges to restrictions on the Appointments Clause, the Supreme Court appears poised to revise the foundations of interbranch delegations of power that undergird the modern administrative state.[53] Mindful that such developments may increase the desirability of a

[52] For an in depth discussion of the purpose and evolution of Congressional agencies, see Kosar, Kevin R., *8 Legislative Branch Support Agencies: What They Are, What They Do, and Their Uneasy Position in Our System of Government*, in CONGRESS OVERWHELMED: THE DECLINE IN CONGRESSIONAL CAPACITY AND PROSPECTS FOR REFORM 128, 128-42 (Timothy M. LaPira, Lee Drutman & Kevin R. Kosar eds., 2020).

[53] In early December 2025, the Supreme Court heard oral arguments in *Trump v. Slaughter* on the constitutionality of restrictions on the President's appointment power. In questioning the Solicitor General, Justice Gorsuch raised the

legislatively-based regulatory entity, we have also included an analysis of the constitutionality of establishing such an entity governing the use of artificial intelligence. That analysis is attached as Appendix A.

1. The Office of Technology Assessment ("OTA")

Established in 1972, the Office of Technological Assessment emerged from a consensus that Congress lacked the expertise to legislate on increasingly complex technical issues.[54] Over two decades, the office produced reports on topics ranging from emerging developments in farming technology to nuclear disarmament. Given the resource-intensive nature of its inquiries, the OTA limited its projects to those initiated by committee chairmen.[55] This feature distinguished the office from the Congressional Research Service, which was and continues to be responsive to rank-and-file legislators. Unlike the OTA, which undertook original, multi-year research projects involving peer review, expert advisory panels, and methodological synthesis, the Congressional Research Service primarily produces rapid-response legal and policy memoranda based on existing sources. CRS work is designed to inform legislators' immediate needs rather than to evaluate emerging technologies through sustained technical inquiry. This methodological difference helps explain why CRS, while indispensable to Congress, has not functioned as a substitute for OTA-style technology assessment.

The office retained only those capabilities necessary to effectuate its research mandate, including the authority to hire research staff, issue subpoenas, and engage contractors.[56] The office

---

question of whether the time had come to reconsider modern administrative rulemaking and "recognize that Congress cannot delegate its legislative authority" to the executive. *See* Transcript of Oral Argument, *Trump v. Slaughter*, No. 25-332 https://www.supremecourt.gov/oral_arguments/argument_transcripts/2025/25-332_7lhn.pdf.

[54] For an in-depth discussion of the history of the OTA, see BRUCE BIMBER, THE POLITICS OF EXPERTISE IN CONGRESS (1996).

[55] 2 U.S.C. § 472(d).

[56] 2 U.S.C. § 473(d).

had no mandate beyond research. The OTA's findings played no statutory or procedurally mandated role in legislation. This relative impotence contrasts with other legislative agencies—like the Congressional Budget Office—whose scoring determinations play a defined role in the legislative process.

Congress structured the OTA to ensure Article I accountability and a non-partisan, expert-oriented approach. The 1972 Act created a bicameral Technology Assessment Board to oversee the agency. The statute also empaneled a Technology Assessment Advisory Council of expert members to be selected by the Board. Council members served in an advisory capacity, offering their high-level technical expertise to the OTA but not engaging with the entity on a full-time basis. Though it varied in size over the years, the OTA employed nearly 130 full-time staff by the early 1990s.

In its 20-year existence, the OTA produced nearly 750 technical reports. Each project required 18 to 24 months to complete.[57] Reports varied widely in scope. Some reports were high-level, surveying developments in an entire field. Others were more responsive to more specific inquiries, including evaluating the computerization of the Social Security Administration and modernization of the FBI's National Crime Information Center.[58]

A. Relevant Expertise

Adopting the OTA model for AI governance issues would equip the Federal government with a deep bench of expertise and help level the information asymmetry between the public and private sectors. Many legislative agencies serve a research and information function, but the OTA model is unique in how it facilitates and channels original, expert research.

---

[57] EMILY BLEVINS, CONG. RSCH. SERV., R46327, THE OFFICE OF TECHNOLOGY ASSESSMENT: HISTORY, AUTHORITY, ISSUES, AND OPTIONS 17 (2020).

[58] Princeton University hosts a comprehensive online catalogue of the OTA's publicly available reports. *See* THE OTA LEGACY, Princeton Univ., https://www.princeton.edu/~ota/.

This capacity flowed in part from the office's structure. The agency's statutory grant provided for a permanent staff, permitted that staff to engage and retain outside scientific expertise, and mandated the inclusion of outside experts on the agency's board.[59] The permanent staff developed deep technical knowledge while the board-level experts provided broader perspective, enabling the OTA to produce both detailed evaluations and high-level frameworks. This dual capacity would particularly benefit AI governance, where Congress needs both overarching regulatory principles and specific assessments of tools deployed in law enforcement.

The original OTA also developed a strong culture of peer review and public-private dialogue that further heightened the quality of its research findings. This approach did not stem from statutory requirements but rather evolved organically over two decades of collaborative research. Over 30 years since the closure of the agency, it will be difficult to reproduce these organic dynamics. However, stakeholders looking to adopt the OTA model for AI should consider amending the original statutory grant to more explicitly codify these cultural norms for a successor agency.

Interviews and historical accounts suggest that OTA's research culture was driven by three reinforcing forces: insulation from day-to-day political demands, sustained engagement with external experts, and long project timelines that rewarded methodological rigor over immediacy. These same features, however, limited OTA's political embeddedness. By restricting its formal relationships primarily to committee leadership and avoiding ongoing engagement with rank-and-file members, OTA failed to cultivate a broad base of institutional champions. In retrospect, the agency's independence and analytic rigor—while central to its credibility—also contributed to its vulnerability when political winds shifted and budgetary protection was needed.

[59] 2 U.S.C. § 472.

B. Policy Influence

The original OTA stood apart from other Federal agencies in the broad respect it commanded among policymakers and legislators.[60] Such "soft power" will be essential to any entity governing AI in criminal justice, a space where well-funded private sector interests will be competing for credibility. In replicating the structure of the OTA, and retaining its name, statutory powers, and resources, some of this soft power may well yet endure in a future entity. These two observations are not in tension but reflect a structural tradeoff inherent in the OTA model. OTA's influence flowed precisely from its lack of formal authority: because it did not propose legislation or issue mandates, its work was perceived as analytically rigorous rather than politically instrumental. At the same time, that same absence of procedural entrenchment meant that OTA's findings could be—and ultimately were—ignored (and the institution itself defunded) without institutional consequence. For AI governance, this suggests that soft power is valuable but insufficient unless paired with mechanisms that embed expert analysis into legislative or administrative workflows.

However, it is important to recognize limitations in replicating this strength. It was arguably the OTA's research culture, and not any element of its statutory powers or mandate, that made the agency uniquely well-respected.[61] Whether such a culture could be replicated is not easily ascertainable. But any adoption of the OTA model should be mindful that the replicable structural elements is only one piece of the puzzle. Channeling the OTA's historic strengths in a new

[60] The OTA's findings were so widely appreciated in Congress that they were considered by eminent legislators as having "raised the level of debate" on issues of science and technology in the legislature. Such respect is sorely needed in the fractious AI policy landscape. *See John Gibbons and Lionel Johns Oral History (Dec. 5, 2006)*, MILLER CTR., UNIV. OF VA., https://millercenter.org/the-presidency/presidential-oral-histories/john-gibbons-and-lionel-johns-oral-history-2006.

[61] BIMBER, *supra* note 6, at 57-60.

organization will require careful stewardship of the entity's culture to ensure it can command the same respect and influence as its predecessor.[62]

C. Organizational Stability

The very elements of the OTA's structure that enabled its strengths put the model's long-term viability in jeopardy. The clearest illustrations of these vulnerabilities come from the 1995 defunding of the original OTA. The OTA's bipartisan design—intended to insulate it from partisan politics—ultimately contributed to its downfall. When Republicans gained control of Congress in 1994 and moved to defund the organization, there was little appetite among Democrats to resist. Similarly, the OTA's statutory grant restricted its contact primarily to the offices of committee chairmen, limiting its relationships with rank and file lawmakers.[63] While this had originally been intended as a way of ensuring the OTA was producing the highest-quality research for the highest-impact initiatives, this feature ultimately alienated the entity when committee leadership suddenly turned over.

There are good reasons to have an expert legislative agency be bipartisan and accountable primarily to committee leadership. But any entity that adopts this model runs the same risk of not having a "champion" within the government to advocate for its continued existence in times of austerity. The OTA's defunding offers a cautionary lesson: bipartisan design is not the same as bipartisan support. An agency can be structured to serve both parties while cultivating deep

[62] In discussing the unique policy influence of the OTA, former Director John Gibbons and former Assistant Director Lionel Johns opined that much of the agency's soft power came from its reputation for viewpoint neutrality, which lent confidence to its outcomes. Such an observation may carry lessons for any future legislative agency handling AI governance issues. *See* Gibbons and Johns Oral History, *supra* note 12.

[63] Nearly every OTA stakeholder interviewed for this white paper highlighted the weak relationships with rank-and-file Congresspersons as a fundamental weakness in the OTA's structure. In his 2024 interview on the OTA, former agency head Peter Blair stressed that any new iteration of the agency would have to address this issue in order to be viable. *See* Santi Ruiz, *How to Assess the Future's Technologies*, STATECRAFT, https://www.statecraft.pub/p/how-to-assess-the-futures-technologies (May 22, 2024) (interview with Peter Blair, former Assistant Director of the Office of Technology Assessment).

relationships with neither. For AI governance, where political valence can shift rapidly—today's consensus on algorithmic fairness may be tomorrow's partisan flashpoint—this institutional rootlessness poses serious risks.

D. Responsiveness

The OTA's in-depth approach to research requires a deliberate and unrushed approach to research that may frustrate stakeholders in the criminal justice space. In its original form, the agency had a turnaround time for research reports of 18 to 24 months.[64] Each month was necessary to the creation of a high-quality product, but such extended lead times posed two serious challenges that would apply with equal, if not greater, force today. The original OTA timeline reflected institutional design choices rather than an inherent feature of rigorous inquiry. A renewed OTA could plausibly accelerate its work through tiered outputs—for example, pairing short-turnaround technical briefs with longer-form assessments, much as GAO's STAA team now does. The principal tradeoff would be depth: faster products would necessarily rely more on existing evidence and expert synthesis rather than original empirical work. For AI in criminal justice, however, such a tradeoff may be acceptable, particularly where delayed guidance risks allowing harmful systems to entrench themselves before oversight occurs.

First, a slow-making research process may dampen the relevance of the agency's findings and its perception among stakeholders. Legislators called on the original OTA to tackle pressing issues like the design of the national power grid or the feasibility of urgently needed defense technology.[65] Such projects can rarely risk delay, and the OTA's deliberate pacing proved immensely frustrating to policymakers. This weakness would continue to haunt a renewed OTA

---

[64] *Id.*

[65] For an in-depth discussion of the trade-offs between the quality of OTA's work product and the timeliness of its reports, *see* AARON FLUITT & ALEXANDRIA GIVENS, IMPROVING TECH EXPERTISE IN CONGRESS, TIME TO REVIVE THE OTA? (2018) https://georgetown.app.box.com/s/2dt0lq0tb6p7kqdf68c7plwewseg3hxs.

in the field of AI in criminal justice. Charged with evaluating AI tools for law enforcement, such an entity would risk seriously diminishing the impact of its recommendations should it not turn out timely reports.

Second, the delay in research would occasionally render the OTA's findings moot, as the technology in question had changed so substantially that the agency's findings were no longer relevant. This issue was less pressing in the 1970s and 1980s. But today, with the rapidly changing pace of technology, such an issue could prove fatal. Artificial Intelligence is advancing at such a pace that 18 to 24 months can be transformative. Taking over a year to evaluate an AI system risks producing recommendations that are outdated upon arrival.

E. Policy Influence

As of now, legislative agencies have limited power to effect binding change.[66] As a research arm of Congress, the original OTA existed far upstream in the legislative process. Its soft power influence notwithstanding, the entity retained virtually no ability to "enforce" its recommendations. The only binding statutory grant of power was the ability to issue subpoenas, a power the OTA never exercised in its 20-year history.[67] In practice, this meant that even the agency's most compelling findings could be ignored without consequence. This stands in stark contrast to executive agencies, which can promulgate binding regulations and impose penalties for noncompliance. For a field like AI in criminal justice—where the stakes include individual liberty and the risk of discriminatory outcomes—relying solely on soft power to shape the deployment of these technologies may prove insufficient.

[66] *See infra* Part III.
[67] *How to Assess the Future's Technologies*, *supra* note 15.

2. The Government Accountability Office ("GAO")

The GAO, as an already-existing agency working on topics related to AI in criminal justice, comes with uniquely low start-up costs. Its proximity to Congress, over 100-year track record, and emphasis on accountability and transparency both to Congress and the public make the agency, in these respects, an excellent candidate for implementing and applying the Task Force's standards.

However, as with the other governance models discussed, the GAO comes with drawbacks. Most notably, the agency is currently experiencing a period of near-unprecedented political instability, facing criticism from the White House and Republicans in Congress as it conducts investigations into President Trump's use of Federal funds. The agency's political vulnerability is amplified by the fact that the most recent Comptroller General, Gene Dodaro, completed his term at the end of December 2025. While Dodaro appointed an Acting Comptroller General, the President, by and with the advice of the Senate, must appoint a new officer to the role in the coming months unless proposals to change the appointments process succeed. In addition, the GAO's research is at times insufficiently detailed, and the agency lacks specialist staff (at least as compared to an OTA-like model), forcing it to resort to external sources of expertise.

The 1921 Budget and Accounting Act established the GAO (called, at the time, the Government Accounting Office) to curtail government spending and reduce Federal debt in the aftermath of the first World War. Though an independent, nonpartisan legislative agency, the GAO reports directly to Congress. Its history is one of successive expansions in its role and capabilities to better facilitate congressional oversight and, through the publication of the GAO's activities on its WatchBlog, accountability to the public. As former Comptroller General Dodaro put it, the GAO's current broad remit involves "looking at whether . . . everything the Federal government

does is [in] accordance with the authorizing legislation of [its] activities."[68] To accomplish this, the GAO audits Federal spending and programs, plays a mandatory role in the Federal budget rescission process, issues legal decisions on bid protests by Federal contractors, drafts policy reports for Congress or to be shared with Federal agencies and the public, and promulgates non-binding recommendations to Federal agencies to which those agencies must, by law, respond, though agencies have discretion as to whether to implement them.

Of particular importance here, in the early 2000s, Congress asked the GAO to experiment with taking up the responsibilities of the OTA, which was by that time already defunct. Over the following two decades, the GAO hired a number of scientists and technology specialists, and in 2018, in recognition of the rapid pace at which technology is advancing and the need for the government to mount an informed response, the GAO established a Science and Technology Analytics Assessment ("STAA") team to research relevant pressing policy issues.

The GAO is led by the Comptroller General, who is appointed by the President by and with the advice of the Senate. As head of a legislative agency, the Comptroller General is not removable by the President—only Congress may do so—thus, once in office, the Comptroller General is largely insulated from the whims of the executive. Second in command to the Comptroller General is the Inspector General, who is neither appointed by the President nor hired with any input from Congress—rather, the Inspector General is an internal hire, which helps to preserve the independence of GAO's leadership. In addition to heads of staff offices, a general counsel, a chief operating officer, a chief financial officer, and other managing employees, the GAO houses 15 "mission teams," which work on particular policy and auditing areas. Potentially relevant to

[68] Gene L. Dodaro & Kevin R. Kosar, *What Does the U.S. Government Accountability Office Do? (with Gene Dodaro)*, AM. ENTER. INST. (Aug. 7, 2023), https://www.aei.org/podcast/what-does-the-u-s-government-accountability-office-do-with-gene-dodaro/.

artificial intelligence in criminal justice are the Defense Capabilities and Management team, Information Technology and Cybersecurity team, Homeland Security and Justice team, and the Science and Technology Analytics Assessment team (STAA), each of which are staffed by policy experts, who are mostly generalists. However, since the creation of STAA in 2018, the GAO has more than tripled the size of that team, including hiring data scientists. GAO also maintains a contract with the National Academies, thus bolstering its ability to convene experts to weigh in on its research and reports.

In its over 100 years of operation, the GAO has generally garnered respect and received robust funding (over $800 million in fiscal year 2024), though it is not immune to allegations of partisanship and related threats to cut funding.

A. Feasibility

The GAO model has the unique strengths that it is currently operational and attuned to issues related to AI in criminal justice. For instance, in September 2023, the GAO issued a recommendation to Immigration and Customs Enforcement (ICE), the Federal Bureau of Investigation (FBI), Customs and Border Protection (CBP), the Department of Homeland Security (DHS), and the Department of Justice (DOJ) on the use of facial recognition services in Federal law enforcement agencies that called for greater training and the creation of policies to protect civil liberties.[69] All five agencies have taken some or all actions recommended by the GAO.[70] Further, the STAA team currently produces both longer reports and shorter two-page technology assessments, which, per the former Comptroller General Gene Dodaro, explain "quickly in simple

[69] U.S. Gov't Accountability Off., GAO-23-105607, *Facial Recognition Services: Federal Law Enforcement Agencies Should Take Actions to Implement Training, and Policies for Civil Liberties*, 42-43 (2023).
[70] U.S. Gov't Accountability Off., *Facial Recognition Services: Federal Law Enforcement Agencies Should Take Actions to Implement Training, and Policies for Civil Liberties* (Sep. 12, 2023), https://www.gao.gov/products/gao-23-105607.

language what a new technology might be, how it's applied, what some of the limitations are, [and] what some of the policy issues are," and are distributed directly to requesting members of Congress.[71] As of August 2023, STAA had produced 38 two-page reports on topics such as quantum computing, advanced batteries, and 5G technology.[72] Topics covered in the longer reports written by STAA and relevant here include AI in healthcare, AI in natural hazard modelling, an assessment of generative AI's environmental and human effects.[73]

These examples of the GAO's current work suggest two natural avenues for implementing the Task Force's standards: through GAO recommendations to Federal agencies, or through the production of policy reports to members of Congress about novel uses of AI in criminal justice.

B. Policy Influence

As a legislative agency, the GAO enjoys an extraordinary degree of proximity to Congress. Employees testify frequently before Congressional committees and members of Congress request reports on topics relevant to their policymaking, thus the GAO is particularly attuned to issues that concern binding rulemakers. The short two-pager reports drafted by STAA on various emerging technologies seem well-suited to informing Congress of the costs and benefits of novel AI tools to be used in the criminal justice system. The effects of such reports, however, would be attenuated, as any legislation drafted in reliance upon them would still need to satisfy the requirements of bicameral passage and presentment.

---

[71] Dodaro & Kosar, *supra*, note 20.
[72] *Id.*
[73] U.S. Gov't Accountability Off., GAO-22-104629, *Technology Assessment: Artificial Intelligence in Health Care* (2022); U.S. Gov't Accountability Off., GAO-24-106213, *Technology Assessment: Artificial Intelligence in Natural Hazard Modeling* (2023); U.S. Gov't Accountability Off., GAO-25-107172, *Technology Assessment: Artificial Intelligence: Generative AI's Environmental and Human Effects* (2025).

Though GAO's recommendations to Federal agencies are non-binding, they have historically garnered considerable respect and are very often implemented.[74] Even if agencies do not implement the GAO's recommendations, they must, by law, respond to the GAO, and indicate how they plan to proceed in light of the Office's recommendation. This is a much more direct form of policy influence and perhaps a more viable avenue by which the Task Force's standards could be implemented.

C. Transparency and Accessibility

The GAO places a premium on transparency and accountability, both key features for effective AI governance. The GAO's staff often testify to the contents of their non-partisan, fact-based reports before Congress. In addition, the Office publishes reports of its research activities and recommendations on its WatchBlog, thus the GAO's work product is accessible both in the sense that it is publicly available online, and that it is generally written in plain, readily understandable language.[75]

D. Relevant Expertise

One concern with the GAO taking the lead on AI in criminal justice is a dearth of subject matter expertise from the staff, who are mostly experts in policy or in financial and program auditing, not in either AI or criminal justice. While STAA is focused on technology, it is led by hard scientists, not individuals with a strong background in AI. Moreover, while the GAO has issued recommendations and drafted some reports related to criminal justice,[76] the GAO does not

[74] In the years from 2019 through 2023, approximately 75-77% of GAO recommendations were implemented by the receiving Federal agency. In 2024, this rate decreased to 70%. U.S. Gov't Accountability Off., GAO-25-900-570, *2024 Performance and Accountability Report*, iv (2025).
[75] See *WatchBlog*, https://www.gao.gov/blog.
[76] *See, e.g.*, U.S. Gov't Accountability Off., GAO-25-107063, *Criminal Investigators: Program-Wide Evaluations and Clear Oversight Responsibilities Could Enhance Training Programs* (2025) ; U.S. Gov't Accountability Off., GAO-24-106348, *U.S. Marshals Service: Actions Needed to Better Identify and Address Detention Condition Concerns* (2024).

have a dedicated domestic criminal justice mission team. Related to the concern about generalist staff, crucial stakeholders in the areas of AI and criminal justice are not located within the GAO, raising the specter that these stakeholders may be insufficiently consulted.

The GAO's ability to leverage external technical expertise mitigates this concern, but only somewhat. The GAO's contract with the National Academies, and the respect commanded by the Comptroller General, enable the GAO to access experts with technical expertise in AI, though these individuals may not be located within the agency itself. For instance, in September 2020, the Comptroller General convened AI experts from across sectors for a Forum on Oversight of Artificial Intelligence. The input of experts notwithstanding, the report produced as a result of this forum takes a relatively high-level view of the issue.[77] This exemplifies the further concern that the GAO's reports are often overly general in nature, which may be linked to the fact that their staff engage primarily in desktop research.[78]

Finally, there is a certain irony in looking to the GAO—an agency whose primary expertise is in saving the government money—to help regulate a notoriously cost-saving and efficiency-increasing technology like artificial intelligence. However, perhaps the fact that the GAO has produced reports like its 2021 AI accountability framework indicates that the agency itself does not see its potential role in AI governance as a contradiction in terms.

E. Organizational Stability and Political Support

The GAO's over 100-year track record and considerable role in the oversight of the use of Federal funds mean that the office is well-established and respected. The GAO's impressive ability

[77] *See* U.S. Gov't Accountability Off., GAO-21-519SP, *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities* (2021).

[78] As Darrell West explained to us in an interview, "My sense is right now, GAO does not have that capability [to conduct in-depth research into AI technology in the criminal justice space], and I don't think they're really set up for the kind of original and in-depth research that we need in the tech area."

to curtail government spending—saving an estimated $70.4 billion in 2023, for instance—illustrates both how effective and deeply entrenched in the Federal government the agency is. However, while the GAO has for much of its history benefitted from bipartisan support, in both the late 1990s and in the past year, the GAO has faced threats from some Republicans in Congress. In the 1990s, these threats were realized, and the GAO was forced to downsize, firing a significant portion of its staff on the heels of a steep budget cut. As researcher Kevin Kosar has noted, this budget cut was both a form of "political punishment" by Republicans in Congress who feared the GAO favored Democrats, and a means of actualizing Newt Gingrich's 1994 campaign, "Contract with America," to shrink the Federal government.[79]

In 2025, similar accusations and threatened budget cuts were levelled against the GAO after the agency investigated President Trump's funding freezes, as it is authorized to do as part of its built-in role in the budget recissions process. The White House Office of Management and Budget (OMB) failed to respond to ordinary GAO requests for information as part of this investigation,[80] and OMB Director Russell Vought publicly derided the agency, suggesting that GAO had played a "partisan role" in its 2019 investigation into President Trump's withholding of aid to Ukraine.[81] Various Republican members of Congress made similar comments suggesting the agency has a partisan bent and could not be relied upon for objective investigations, despite the fact that both Republicans and Democrats in Congress rely heavily on the GAO's work.[82] In

[79] Kosar, *supra* note 1 at 138.
[80] Katherine Tully-McManus, *Trump administration ignores watchdogs amid 39 funding investigations, GAO says*, POLITICO (Apr. 29, 2025), https://www.politico.com/live-updates/2025/04/29/congress/trump-administration-ignores-watchdogs-amid-39-funding-investigations-gao-says-00315459.
[81] Jennifer Scholtes, Jordain Carney & Katherine Tully-McManus, *GOP declares war on GAO*, POLITICO (May 27, 2025), https://www.politico.com/news/2025/05/27/gop-declares-war-on-gao-00369762?ICID=ref_fark&utm_content=link&utm_medium=website&utm_source=fark.
[82] *Id.*

line with this, one Republican budget plan proposed last year included a near 50% cut to the GAO's budget (ultimately unsuccessful).

Though the GAO escaped such steep budget cuts, it remains enmeshed in a political thicket: the GAO's investigations into President Trump's use of Federal funds are ongoing, and Gene Dodaro, Comptroller General for the past 15 years, completed his term on December 30, 2025. While Dodaro appointed Orice Williams Brown as Acting Comptroller General, a replacement must be selected by the President by and with the advice of the Senate. More specifically, a bipartisan and bicameral committee of 10 members of Congress must recommend a shortlist of at least three candidates to the President, who in turn may request additional options or choose one to be recommended for confirmation by the Senate. In recognition of the importance of the role and the risks to a successful appointment posed by the current political climate, a Democrat in the House of Representatives recently proposed a bill modifying the Comptroller General's appointments process to authorize "Congress or the President, respectively," to choose from among the shortlisted candidates, rather than leaving this responsibility with the President alone.[83] While similar modifications have also been proposed by a former Comptroller General, David Walker, it is unclear whether the bill will gain any traction.[84] With a term length of 15 years for the next Comptroller General, the selection (if, in fact, someone is successfully confirmed) in the coming months will be critical for the GAO's future.

Once in office, however, the Comptroller General is only removable by Congress through impeachment or a joint resolution, thereby insulating them from the Executive Branch. The

[83] H.R. 6517, 119th Cong. (2025).

[84] *See* David Walker, *Choosing the next comptroller general is too important– don't play politics*, THE HILL (June 4, 2025), https://thehill.com/opinion/congress-blog/5330724-choosing-the-next-comptroller-general-is-too-important-for-playing-politics/.

Inspector General is further insulated from both the President and Congress, as they are internally hired by the GAO.

C. Sentencing Commission Model

This section explores how the institutional design features of sentencing commissions can inform a governance framework for overseeing the use of artificial intelligence in the Federal criminal justice system. By comparing sentencing commissions at the federal and state levels, our findings explore lessons from their institutional design features that balance independence, expertise, and transparency.

Specifically, this section evaluates the Federal U.S. Sentencing Commission and two state commissions: North Carolina and Minnesota. Across jurisdictions, sentencing commissions demonstrate that independent expert bodies, when grounded in transparent procedures and guided by empirical evidence, can effectively translate complex insights into actionable and legitimate policy goals. This approach is especially relevant for AI governance in public systems where similar challenges arise in ensuring fairness, accountability, and adaptability.

Our key findings include the following: First, an independent hybrid structure outside the Executive Branch that performs legislative functions can address questions of accountability and autonomy. However, such entities should include a robust research staff along with subject-matter decision-makers in order to draw on the necessary technical expertise. Second, data-driven guidelines can bring empirical competency, technical nuance, and legitimacy, provided they are transparent and driven by iterative review. For instance, the U.S. Sentencing Commission solicits public comment and holds public hearings. Third, the state-level models reveal how administrative flexibility and local participation can coexist, suggesting that AI governance can benefit from multi-level coordination and assessment rather than a purely Federal solution.

By "administrative flexibility," we do not mean ad hoc discretion, but rather the capacity of state-level commissions to adapt institutional design, staffing practices, and analytic priorities to local legal structures and data environments. Both the North Carolina and Minnesota commissions illustrate how standing bodies embedded in state systems can effectively recalibrate research agendas, respond to legislative changes, and integrate feedback from courts and practitioners. "Local participation" refers not only to geographic proximity, but to the incorporation of decision-makers drawn directly from the institutions and communities governed by the commission's outputs. Together, these features suggest that state-level entities can serve as effective sites of experimentation and learning, supporting a multi-level AI governance architecture rather than a single centralized federal solution.

Each of the key findings identified above is contingent on the institutional design attributes that structure this analysis. The viability of an independent hybrid entity, potentially outside the Executive Branch, for example, depends critically on the presence of relevant expertise—both among decision-makers and within a standing technical staff—without which independence risks becoming insulation rather than accountability. Likewise, the ability of data-driven guidelines to generate legitimacy and technical nuance is not only a function of transparency and iterative review, but also of organizational stability, which allows empirical monitoring, revision, and trust to accumulate over time. Making these dependencies explicit clarifies that institutional form alone is insufficient; performance turns on how specific design attributes are combined.

Case Studies

1. U.S. Sentencing Commission

The United States Sentencing Commission is an independent agency in the Judicial Branch charged with developing the Federal sentencing guidelines. Created by Congress through the

Sentencing Reform Act of 1984, the Commission occupies a constitutionally unusual position: although housed in the Judicial Branch, it exercises quasi-legislative and quasi-executive powers. Its hybrid structure was upheld by the Supreme Court in *Mistretta v. United States* (1989),[85] which recognized the Commission as a "special" entity. The Commission offers several instructive lessons for the design of a new body tasked with governing AI in federal criminal justice.

A. Relevant Expertise

The Commission consists of seven commissioners appointed by the President and confirmed by the Senate, no more than four of whom may belong to the same political party. Beyond the commissioners themselves, the Commission has up to 100 professional staff devoted to research and data analysis. For AI governance, this model highlights the importance of pairing (bipartisan) subject-matter decision-makers with a robust technical and analytical staff. However, the Commission's expertise is tightly tethered to sentencing policy and judicial decision-making. As one interviewee noted, the Commission is "one or two steps removed" from operational practices such as law enforcement or procurement. Thus, an AI governance entity would likely require a broader mix of expertise that goes beyond the Commission's court-centered orientation.

B. Responsiveness

One of the Commission's most notable strengths is its procedural flexibility. Because it is exempt from most administrative law statutes that govern Executive Branch agencies, the Commission can move relatively quickly. It proposes guideline amendments on an annual cycle, incorporates public comment, and submits finalized amendments to Congress by a statutory deadline. For AI governance, this level of nimbleness is attractive. AI technologies evolve rapidly,

[85] 488 U.S. 361 (1989).

and an oversight body that cannot adapt in near-real-time risks irrelevance. At the same time, one interviewee noted that the Commission has historically struggled to implement its statutory "monitoring" mandate. This suggests that exemption from administrative burdens must be paired with sufficient staffing if an AI-focused entity is to keep pace with technological change.

C. Transparency and Accessibility

Although exempt from transparency statutes such as FOIA and the Sunshine Act, the Commission voluntarily operates with a high degree of transparency. It solicits public comments, holds hearings, and releases comprehensive sentencing datasets. While openness is not legally required, it is indeed embedded in the Commission's institutional culture and legislative history. This model offers a lesson for AI governance: formal exemptions from administrative law need not equate to opacity. An AI oversight entity could also combine procedural flexibility with self-imposed transparency norms including public reporting, open datasets, and participatory processes. However, reliance on voluntary practices also poses risks. For instance, the Commission used to employ an informal "open door" policy allowing any interested party to sit in on their meetings. But as their discontinuation of the policy demonstrates, transparency mechanisms that are not legally entrenched can be abandoned over time. For AI governance—where trust deficits are already acute—transparency obligations may need firmer statutory grounding.

D. Organizational Stability

The Commission's structure promotes long-term stability through staggered six-year terms for commissioners and the Chair's central role, particularly over budgetary decisions. Together, these design choices have allowed the Commission to maintain a consistent institutional identity across administrations. For an AI governance body, such stability would be a significant asset.

Indeed, different administrations have already exhibited vastly different AI goals, with the Biden administration emphasizing transparency and algorithmic fairness and the Trump Administration seeking deregulation and accelerated development.[86] At the same time, the Commission's experience also illustrates a tradeoff: stability can turn into inertia. Vacancies and resource constraints have at times limited the Commission's effectiveness. A future AI entity would need mechanisms to preserve continuity without allowing leadership gaps to undermine its mission.

E. Startup Feasibility and Policy Influence

Despite its many strengths, the Commission's model is difficult to replicate. The Commission was placed in the judiciary because people saw sentencing as a core judicial function; in contrast, an entity intended to guide AI use in criminal justice would face significant pushback if housed in the judiciary. At the same time, the Commission demonstrates how an expert body can exert substantial policy influence without direct enforcement power. Its guidelines shape sentencing nationwide, and Congress has at times explicitly relied on the Commission to advance reforms that stalled legislatively. Thus, an AI governance body issuing nonbinding guidance could still meaningfully shape practice.

2. North Carolina Sentencing and Policy Advisory Commission

The North Carolina Sentencing and Policy Advisory Commission's enabling legislation[87] authorizes the commission to function as an independent and permanent body, which is housed in

---

[86] *Compare* Removing Barriers to American Leadership in Artificial Intelligence, White House (Jan. 23, 2025), https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/, *with* Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, White House (Oct. 30, 2023), https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.

[87] North Carolina General Statutes § 164-35 to -50 (2025).

the Judicial Branch. It was established by the General Assembly in 1990 for reasons which strongly mirror the need for a regulatory body for AI in criminal justice: a lack of uniformity and predictability amongst sentencing practices compelled the need for an independent body to revise and rewrite sentencing laws.[88] The Commission was subsequently enacted as a permanent body in 1996 to reify its role in monitoring the criminal justice system more broadly. There are several institutional design features of the North Carolina Sentencing and Advisory Commission which are foundational to its reputable tradition, longstanding credibility, and seamless integration within the fabric of the criminal justice system. While there are startup and policy influence limitations, the Commission's strengths in relevant expertise, responsiveness, transparency/accessibility, and organizational stability ultimately provide a template for AI governance in criminal justice.

A. Relevant Expertise

North Carolina is exemplary among state sentencing commissions for its breadth and depth in criminal justice expertise. Its 29-commissioner membership structure is statutorily authorized and includes judges, prosecutors, public defenders, formerly incarcerated individuals, community members, and a diverse array of individuals selected to represent the criminal justice system.[89] Members are appointed by the Governor, the Chief Justice, the Speaker of the House, and the President Pro Tempore of the Senate. For AI governance, it is critical that decisions are informed by the judgement of experts across the criminal justice system as powerfully demonstrated by this commission. This expansive diversity of criminal justice expertise and viewpoints is additionally supported by designated staff, detailed in the following section, who supply technical expertise.

---

[88] Lorrin Freeman, *The North Carolina Sentencing and Policy Advisory Commission: A History of its Creation and its Development of Structured Sentencing* (updated Aug. 2014), North Carolina Judicial Branch.
[89] North Carolina General Statutes § 164-35 to -50 (2025).

B. Responsiveness

Through the appointment of staff members, the commission is designed to produce timely and responsive work products. The staff, led by an Executive Director, is composed of a research and policy team: one Associate Director for Policy/Staff Attorney, one Associate Director for Research, one Administrative Secretary, one Senior Research and Policy Associate, and five Research and Policy Associates. The research team supports the work of the Commission by acquiring and compiling data for the Commission to conduct studies and produce reports on their findings; the policy team tracks updates within the General Assembly and responds to legal and policy questions. For AI governance, this organizational structure would allow for the responsiveness necessary for tracking rapid advancements in AI instruments as procured within criminal justice.

C. Transparency and Accessibility

The Commission has a longstanding tradition of transparency and accessibility, which will be similarly indispensable for AI governance. Meeting minutes, agendas, and study reports are publicly accessible through the North Carolina Sentencing and Policy Advisory Commission website.[90] The statute itself demands, "all meetings of the commission shall be open to the public and the information presented to the commission shall be available to any state agency or member of the General Assembly."[91] For AI governance, it is similarly imperative that work products and processes are as accessible as possible.

[90] *North Carolina Judicial Branch*, North Carolina Courts, https://www.nccourts.gov/ (last visited Dec. 22, 2025).
[91] North Carolina General Statutes § 164-35 to -50 (2025).

D. Organizational Stability

The Commission has operated successfully for 35 years due to its stability in leadership. Notably, the Chairman, a sitting or former justice or judge who is appointed by the Chief Justice of the North Carolina Supreme Court, is an integral figure in ensuring this longevity. Given the Commission's diversity in representation, disagreements among members are inevitable and may lead to challenges in policy alignment. Effective leadership within the Commission is thus essential in ensuring that all perspectives, especially dissenting minorities, are judiciously considered with fair and equal weight. A commission governing the use of AI may similarly benefit from the leadership of a judge who may effectively mediate disagreements in order to ensure stability over time.

E. Startup Feasibility and Policy Influence Limitations

There are a few limitations that must be considered in the context of the Commission's strengths. Because the North Carolina Sentencing and Policy Advisory Commission was established by statute, the legislative process required to establish a similar commission for AI governance may demand some time and political will. Second, while the Executive Director may draft recommendations based on the Commission's findings, these recommendations must be approved by the legislature and signed by the governor. However, it is important to consider that these limitations are a necessary compromise to provide for a level of political independence. Additionally, the power of nonbinding recommendations should not be understated, as influence may still be exerted if developed by a well-respected institution.

3. Minnesota Sentencing Guidelines Commission (MSGC)

Minnesota was the first state to establish a sentencing commission which has become an influential model for structured and empirically grounded sentencing reforms.[92] Created by statute in 1978 and operational since 1980, the Minnesota Sentencing Guidelines Commission (MSGC) is a relatively small body with a unique hybrid structure, an independent agency intended to balance judicial expertise, legislative accountability, and executive implementation.

A. Relevant Expertise

The MSGC's design institutionalizes broad criminal justice expertise through a deliberately pluralistic membership structure. The Commission consists of thirteen members representing judges, prosecutors, defense counsel, law enforcement, probation, corrections, rehabilitative services, academia, and the public. The Commission was designed to promote proportionality, uniformity, and transparency in sentencing through binding guidelines informed by continuous data collection and analysis.[93]

Statutory requirements ensure that one public member represents crime victims or victim advocates and another is a formerly incarcerated individual who has completed their sentence. Members are appointed through a distributed process involving both the governor and judicial actors, reinforcing cross-branch legitimacy.

B. Responsiveness

The MSGC operates through a continuous, iterative governance model rather than fixed reform cycles. The law requires the Commission to meet regularly, typically monthly, to review

[92] 1978 Minn. Laws ch. 723, art. 1, § 9; see Minn. Stat. § 244.09, subd. 1 (2019).
[93] Michael Tonry, *The Untold Story of America's First Sentencing Commission*, 21 Fed. Sentencing R. 265 (2009).

sentencing data, consider research findings, and evaluate whether guideline modifications are warranted. There is no rigid schedule for reform; its updates are driven by evidence and observed practice. Commission staff described how elevated judicial departure rates are treated not as judicial error but as diagnostic signals that guidelines may be misaligned with on-the-ground realities. These signals prompt Commission discussion and, where appropriate, targeted research or guideline revisions. At the same time, responsiveness is constrained by limited data infrastructure. Sentencing data are drawn from multiple non-interoperable systems operated by courts, probation agencies, and corrections. Staff must manually reconcile discrepancies, validate records, and in some cases contact judges directly to clarify sentencing information. While this ensures data accuracy, it significantly slows analytic turnaround.

C. Transparency and Accessibility

All MSGC's meetings are open to the public; proposed guideline amendments are subject to public hearings and comment; and Commission materials, including annual reports and detailed statistical appendices are published online.

Commission staff emphasized that transparency is essential to maintaining legitimacy, given the statutory requirement and binding nature of the guidelines. Chairs often pursue consensus or near-unanimous votes before advancing proposals, recognizing that broad internal agreement facilitates multi-stakeholder acceptance and legislative review. These consensus-oriented practices were repeatedly cited as central to the Commission's credibility with courts, legislators, and the public.

D. Organizational Stability

The MSGC's nearly five decades of continuous operation reflect a high degree of organizational stability. This durability is rooted in its statutory authorization, staggered appointments, and distributed appointment structure. Commissioners serve staggered terms, preventing wholesale turnover with changes in political leadership. The governor's appointment of the Chair is balanced by judicial appointments and statutory membership requirements.

E. Startup Feasibility and Policy Influence Limitations

The MSGC model benefits from being embedded within the criminal justice system. Its placement in the Judicial Branch has allowed it to remain politically durable and largely insulated from partisan swings. Commission staff noted that this institutional neutrality has helped the body during periods of broader criminal justice politicization. However, the model also has limits. While the Commission's guidelines are legally binding on courts, its authority does not extend to law enforcement agencies, procurement decisions, or technology vendors. Its influence outside sentencing operates primarily through data, analysis, and agenda-setting rather than direct regulatory control. Moreover, the Commission's small size limits its ability to absorb significantly expanded responsibilities without additional legislative support.

Conclusion

Taken together, the three commissions illustrate that the balance between subject-matter decision-makers and technical staff is not incidental but outcome-determinative. The Federal Commission's relatively small, professionally homogeneous membership is offset by a large, highly specialized research staff capable of sustained empirical analysis, whereas North Carolina's unusually broad and diverse commission membership compensates for a much smaller staff by

embedding experiential and institutional knowledge directly into deliberation. Minnesota occupies a middle position, pairing pluralistic membership with a lean staff that is effective within its narrow sentencing mandate but strained by data and infrastructure constraints. For purposes of AI governance, these contrasts suggest that no single configuration is optimal in the abstract: commissions with narrower mandates can rely more heavily on staff expertise, while bodies tasked with evaluating sociotechnical systems across the criminal justice pipeline likely require both broader decision-maker representation and substantially more robust analytic and data infrastructure.

D. Congressional Committee Model

Joint Congressional Committees (JCCs), such as the enduring Joint Committee on Taxation, highlight a non-legislative governance model that centralizes bipartisan technical expertise for cross-jurisdictional issues like taxation, producing rigorous analyses, revenue estimates, and recommendations without direct bill-reporting authority. Similarly, the House and Senate Judiciary Committees serve as primary hubs for criminal law and civil rights oversight, leveraging subcommittees, hearings, and legislative drafting to address policing, sentencing, and adjudication, though they often rely on external experts amid limited in-house technical staff.

These congressional committee models offer distinct strengths but face notable limitations for AI governance. JCCs excel in organizational stability, relevant expertise through full-time non-partisan staff, and indirect policy influence via authoritative reports that unify bicameral perspectives, yet they score lower on startup feasibility, requiring concurrent resolutions or statutes. They also score low on responsiveness, due their deliberative workflows. Judiciary subcommittees, by contrast, provide high feasibility via simple internal votes, strong transparency

through public hearings, and direct legislative power, but suffer from fragile stability tied to partisan shifts and inconsistent technical depth.

The suitability of congressional committees to advance AI governance in criminal justice hinges on their complementary roles. JCCs could form a stable technical backbone for long-term analysis and coordination across committees like Judiciary, Commerce, and Appropriations, while Judiciary subcommittees deliver agile, enforceable responses through hearings and bills. Together, they address Congress's institutional gaps in expertise and cross-jurisdictional capacity, but neither alone suffices; a hybrid approach, pairing a JCC's endurance with subcommittee speed, would better sustain oversight amid AI's rapid evolution, though political polarization remains a barrier to implementation.

The analysis concludes that while a joint congressional committee offers broad, long-term coordination, establishing a dedicated Judiciary Subcommittee on AI Governance represents a more feasible and legislatively focused approach in the short term.

Case Studies

This section consists of a detailed comparative institutional analysis of the two congressional committee models—the JCC and the judiciary committees—to assess potential governance models. It also presents criteria to measure the suitability of AI for the criminal justice system.

1. Joint Congressional Committees (JCC)

Joint Congressional Committees (JCCs) are permanent or temporary panels composed of members from both the House and the Senate. Unlike standing committees, which hold the power to report or propose legislation to the floor, JCCs are non-legislative. Joint congressional committees have historically been established for cross-jurisdictional, technically complex issues

(such as taxation), where a bicameral, centralized resource would enhance Congress's capacity. The Joint Committee on Taxation (JCT) is one of the most successful JCCs and is currently still active.

Joint congressional committees are typically established as investigative or advisory bodies created to study broad issues of shared concern, such as economic policies[94] or atomic energy.[95] JCCs do not possess direct legislative authority, meaning they cannot report bills to the floor for a vote.[96] Their influence is instead derived from their capacity to conduct high-level oversight and investigations, produce rigorous policy analyses and reports, and, most importantly, facilitate coordination between the two chambers, thereby helping to unify congressional perspectives on complex, often nationally relevant issues.

The Joint Committee on Taxation (JCT) is a well-staffed, neutral, independent, and widely respected participant in the development of tax legislation.[97] The JCT is tasked with investigating the operation, effects, and administration of internal revenue taxes, as well as exploring methods for simplifying them. It reports its findings and recommendations to the House Committee on Ways and Means, the Senate Committee on Finance, or directly to Congress. The JCT also reviews certain proposed tax refunds or credits.[98]

To carry out these responsibilities, the JCT is authorized to obtain and inspect tax returns and related information, hold hearings, compel witness testimony, procure printing services, and make necessary expenditures. Additionally, the JCT may, with proper approval, obtain tax data

---

[94] Joint Econ. Comm., U.S. Congress, *About* , https://www.jec.senate.gov/public/index.cfm/about.
[95] Stanford Univ. Libraries, *Perspectives on Nuclear Technology*, https://exhibits.stanford.edu/atomic-energy.
[96] U.S. Senate, *Frequently Asked Questions About Committees*, https://www.senate.gov/committees/committees_faq.htm.
[97] Donald L. Korb, James J. Erney, Gregory J. Gawlik, Rethinking Refund Review: Understanding the Joint Committee on Taxation, 4 Corp. Bus. Tax'n Monthly 3 (2002).
[98] Joint Comm. on Taxation, U.S. Congress, *Statutory Basis*, https://www.jct.gov/about-us/statutory-basis/.

directly from the IRS or other executive agencies to support its investigations, reports, and studies.[99]

JCT's creation reflected the recognition that tax policy required specialized, non-partisan technical staff beyond what individual committees could sustain. Similar rationales now justify a JCC for AI governance: AI is cross-cutting (touches Judiciary, Commerce, Intelligence, Homeland Security, Appropriations), technically intricate, and subject to rapid evolution. A JCC would centralize expertise, produce durable technical analyses, and coordinate bicameral approaches.

A JCC can be established in two ways: (i) passing of a concurrent resolution by both the House of Representatives and the Senate; or (ii) establishment of a joint committee by law (e.g., the Joint Committee on Taxation).[100] Membership and appointment procedures vary as per the establishing document, but in general JCCs draw senior members from both chambers and often alternate chairmanships between the House and Senate. For instance, the Joint Committee on Taxation alternates its chair between the Chair of the Senate Finance Committee and Chair of the House Ways and Means Committee.[101] JCCs typically maintain a bipartisan staff, which includes subject-matter experts. The Joint Committee on Taxation staff, for example, includes economic and tax policy experts who provide technical analyses for the committee's investigations and reports.[102]

The work products of JCCs vary by their mandate but generally include investigative reports, policy analyses, studies, and recommendations to Congress. The Joint Committee on

---

[99] *Id.*

[100] Revenue Act of 1926, ch. 27, § 1203, 44 Stat. 9, 129 (establishing the Joint Committee on Taxation to investigate the practices of the Bureau of Internal Revenue).

[101] Cong. Rsch. Serv., R46786, Rules Governing House Committee and Subcommittee Assignment Procedures (2025), https://www.congress.gov/crs-product/R46786.

[102] Official Congressional Directory, at 710 (2d Sess. 2024), https://www.govinfo.gov/content/pkg/CDIR-2024-04-25/pdf/CDIR-2024-04-25.pdf.

Taxation produces detailed reports on the operation, effects, and administration of internal revenue taxes, as well as revenue estimates and analyses of proposed tax legislation. Other JCCs may produce oversight reports, summaries of inter-chamber studies, or recommendations on specific legislative matters.

Congress possesses powerful institutional tools such as appropriations, oversight, investigatory subpoena power, and lawmaking authority, but its existing research support agencies such as the Congressional Research Services, provide primarily short-term and legal or policy analysis rather than continuous specialized technical oversight. Congress lacks sustained technical capacity and a nimble cross-committee structure to manage rapidly evolving AI in the criminal justice system. The JCC model could directly target this institutional gap by prioritizing enduring technical expertise, fostering advisory influence, and enabling cross-committee collaboration, although it would trade out startup feasibility, direct enforceability, and speed of response.

A. Startup Feasibility

The political and procedural likelihood of establishing a JCC on AI in Criminal Justice is relatively unlikely. The proposal to establish one envisions the creation of a new, non-partisan, full-time congressional entity, which would likely require substantial negotiation over structure, staffing, and jurisdiction, and would necessitate formal authorization through concurrent resolution or statute. These procedural demands make formation more complex than reactivating or expanding an existing committee. While dependence on bipartisan leadership can pose an institutional hurdle in a politically polarized environment, recent expressions of concerns surrounding AI's societal and security implications suggest that this issue could also potentially attract cross-party interest. Therefore, the political and procedural feasibility of establishing a JCC

is uncertain as it would hinge on whether or not emerging AI oversight efforts will be a governmental priority.

B. Relevant Expertise

The JCC model is explicitly designed to remedy Congress's current lack of sustained technical staff capable of evaluating AI models and methodologies in the criminal justice context. It would serve as a non-partisan, full-time technical resource, ensuring that government bodies have access to reliable, independent expertise on AI governance issues. The JCC would embed deep technical and methodological competence within the legislative process.

C. Transparency and Accessibility

The JCC would promote transparency primarily through its public-facing analytic output rather than through direct legislative proceedings. Periodic public reports, scored cost/benefit analyses, and published recommendations are grounded in public transparency. A JCC would create a publicly accessible record of its technical assessments and governance guidance. As a central repository for AI governance information across relevant committees, it would enhance informational accessibility for legislators, staff, and external stakeholders, even if it is not itself a primary forum for public testimony.

D. Organizational Stability

Institutional durability is a core strength of the JCC model. Envisioned as a permanent, non-partisan entity with full-time staff, it would develop long-term institutional memory on AI and criminal justice that transcends individual election cycles. By facilitating cross-chamber coordination and serving as a standing technical reference point for multiple authorizing and appropriations committees, it would provide a stable foundation for complex policy analysis and incremental standard-setting over time.

E. Policy Influence

Although primarily advisory and lacking direct bill-reporting authority, the JCC could wield significant indirect policy influence. Its role is analogous to that of the Joint Committee on Taxation, which regularly shapes legislative outcomes in taxation policies, at least at the technical level. Authoritative technical assessments, cost/benefit analyses, and governance recommendations from a JCC on AI could meaningfully guide multiple committees' legislative drafting, oversight priorities, and funding decisions, even in the absence of formal lawmaking power.

F. Responsiveness

The JCC's emphasis on research, methodological documentation, and non-partisan credibility may limit its responsiveness to rapidly evolving legislative demands. Producing high-quality technical analysis and scored evaluations is inherently time-consuming, creating a risk that its outputs lag behind urgent political or oversight needs in a fast-moving technological and polarized environment. For instance, as part of the Joint Committee on Taxation's responsibility, its analyses predict Federal revenues under a proposal with those expected under current law. The JCT has lengthy processes for producing these revenue estimates, which typically require integrating IRS data with various economic models, incorporating behavioral responses, undergoing interdisciplinary review, etc. This deliberative model (which is essential for credibility) creates risks where outputs lag behind urgent political or oversight needs, especially in the context of fast-changing technologies such as AI. In the area of AI and criminal justice, where legislative and ethical questions can shift rapidly with technological breakthroughs or crises, maintaining both speed and analytical integrity would be a persistent institutional challenge.

Across the identified criteria, the JCC model offers strong performance on relevant expertise, organizational stability, and indirect policy influence, with moderate strengths in transparency and accessibility through its public reports and shared analyses. However, it scores lower on startup feasibility due to the political and procedural hurdles of creating a new joint institution, and on responsiveness and direct enforceability, given its advisory mandate and methodologically intensive workflow. These limitations suggest that while the JCC could become the backbone of Congress's long-term technical capacity on AI in criminal justice, it may need to operate in concert with more agile, enforcement-oriented mechanisms to fully meet emerging governance needs.

2. Judiciary Committees

The Judiciary Committees in the House and Senate serve as the primary venues for legislation and oversight relating to criminal law, civil rights, constitutional protections, and the Federal courts.[103] Unlike more specialized or sectoral committees, they possess broad jurisdiction over core questions of law such as due process, discrimination, surveillance, and procedural fairness that are directly implicated by the deployment of AI in the criminal justice system. This combination of wide-ranging subject-matter authority and longstanding institutional experience makes the Judiciary Committees natural hubs for examining how AI tools affect investigations, adjudication, sentencing, and correctional practices, and for developing baseline governance standards in response.

The existence of committees dedicated to justice and judicial matters dates back to the early years of the Republic. The Senate Judiciary Committee, established in 1816,[104] and the House

[103] *About the Committee*, U.S. Senate Comm. on the Judiciary, https://www.judiciary.senate.gov/about/committee; *About*, H. Comm. on the Judiciary, https://judiciary.house.gov/about.
[104] *About the Committee*, U.S. Senate Comm. on the Judiciary, https://www.judiciary.senate.gov/about/committee.

Judiciary Committee, established in 1813,[105] are among the oldest and most influential standing committees in Congress. They were created out of the necessity for a specialized body to handle the complex and technical work of drafting and scrutinizing laws that define Federal crimes, civil rights, constitutional amendments, and the structure of the Federal courts.

A standing committee is a permanent legislative panel established by the House and Senate rules that has jurisdiction over a particular subject area. Unlike joint congressional committees, which primarily serve an advisory or investigative function without direct lawmaking power, standing committees possess the crucial authority to report bills and proposed amendments to their respective chambers for a vote, making them indispensable to the legislative process.

The House and Senate Committees on the Judiciary serve as the primary legislative and oversight engines for the American legal system and the protection of constitutional order. Unlike advisory bodies, these standing committees possess the direct authority to draft, mark up, and report legislation to their respective chambers.[106] Their mandate is broad, encompassing the development of criminal law, the regulation of immigration and naturalization, and the advancement of civil liberties.[107] Furthermore, they serve as the primary gatekeepers for laws governing emerging technologies including intellectual property law governing patents, copyrights, and trademarks.

The House and Senate Judiciary Committees are structured as standing committees, with membership generally reflecting the partisan balance of their respective chambers. These committees are established by the rules of the House and the Senate at the beginning of each new Congress.

---

[105] *About*, H. Comm. on the Judiciary, https://judiciary.house.gov/about.

[106] *Committee Consideration*, U.S. Congress, https://www.congress.gov/legislative-process/committee-consideration.

[107] *About the Committee*, U.S. Senate Comm. on the Judiciary, https://www.judiciary.senate.gov/about/committee.

Members are appointed by the chamber's leadership, usually upon recommendation from the party caucuses. Due to the high-profile nature of their work, they often attract members with legal backgrounds or strong interests in constitutional and legal issues. The chairperson is typically a senior member of the majority party, and the ranking member is a senior member of the minority party.

In order to handle their expansive jurisdiction, these committees are divided into specialized subcommittees (e.g., on Courts, Antitrust, Immigration) that focus on specific areas of law and policy. These committees maintain robust professional staff, often including legal experts and policy analysts, who provide legislative research, draft bills and reports, and prepare for hearings and markups.

The core output of these committees is legislation, involving the drafting, rigorous amendments, and reporting of bills that define national policy on various legal matters. These legislative efforts are often accompanied by committee reports, which serve as the official record of a bill's intent and provide critical guidance for judicial interpretation.

Beyond written products, the committees utilize hearings as essential public forums for debating the most pressing legal and constitutional issues of the day. Recent sessions have seen Judiciary subcommittees tackle the complexities of algorithmic bias, facial recognition, and AI's implications for intellectual property and due process.[108] While these hearings demonstrated the committees' formidable power to summon agency testimony and catalyze legislative text, they also

[108] The Good, the Bad, and the Ugly: AI-Generated Deepfakes in 2025: Hearing Before the S. Comm. on the Judiciary, 119th Cong. (2025), https://www.judiciary.senate.gov/committee-activity/hearings/the-good-the-bad-and-the-ugly-ai-generated-deepfakes-in-2025; Examining the Harm of AI Chatbots: Hearing Before the S. Comm. on the Judiciary, 119th Cong. (2025), https://www.judiciary.senate.gov/committee-activity/hearings/examining-the-harm-of-ai-chatbots; Protecting Our Edge: Trade Secrets and the Global AI Arms Race: Hearing Before the H. Comm. on the Judiciary, 119th Cong. (2025), https://judiciary.house.gov/committee-activity/hearings/protecting-our-edge-trade-secrets-and-global-ai-arms-race; AI Crossroads: Nationwide Strategy or Californication?: Hearing Before the H. Comm. on the Judiciary, 119th Cong. (2025), https://judiciary.house.gov/committee-activity/hearings/ai-crossroads-nationwide-strategy-or-californication.

exposed significant knowledge constraints. Much of the testimony relied heavily on academia or industry experts,[109] highlighting a critical gap where technical expertise may be influenced by incentives that do not necessarily align with the public interest.

A. Startup Feasibility

It is highly feasible to establish a dedicated Judiciary Subcommittee on AI and the criminal justice system. The subcommittee could be created through a simple majority vote within the Judiciary Committee, without requiring new statutory authority, making it a procedurally efficient option even in a divided Congress. By leveraging existing committee structures, this model minimizes institutional startup costs while enabling rapid initiation of oversight and legislative activity.

B. Relevant Expertise

The Judiciary Committee's jurisdiction over criminal law, civil rights, and the Federal courts positions it as the natural institutional home for AI governance within the criminal justice system. This existing mandate ensures substantial legal and policy expertise on issues such as due process, discrimination, and procedural fairness, which are central to AI deployment in policing, sentencing, and corrections. However, technical AI expertise among staff is likely to be limited, as hiring patterns traditionally favor legal generalists and legislative staffers rather than engineers or data scientists. Given the committees' already burdened agendas, relying solely on existing staff to handle complex AI issues on an ad hoc basis is not likely to provide the sustained, specialized attention that these tools require. Addressing this gap would therefore require dedicated, full-time AI and data governance staff within the Judiciary Committees, or at minimum, an institutionalized technical advisory capacity which is supported by the budgetary commitments and political

[109] *Id.*

backing to ensure the issue of AI and criminal justice does not lose out to competing legislative priorities.

C. Transparency and Accessibility

The subcommittee model offers strong transparency and stakeholder accessibility. Public hearings, published transcripts, and written testimony would provide visible forums for civil society, impacted communities, technologists, and practitioners to shape the AI governance agenda. The subcommittee's authority to hold hearings on agency AI usage, receive evidence on procurement and performance, and issue reports and white papers would further institutionalize accessible, public-facing deliberation on AI in criminal justice.

D. Organizational Stability

Organizational stability represents a relative weakness of the subcommittee model. As a unit internal to the Judiciary Committee, the subcommittee's existence, jurisdictional scope, and agenda remain subject to internal committee rules and decisions at the start of each Congress. Unlike independent commissions or entities established by statute, a subcommittee can be restructured, downgraded, or dissolved with changes in partisan control or leadership priorities, reducing continuity for long-term AI governance initiatives.

E. Policy Influence

The proposed subcommittee would wield substantial policy influence over AI governance in the criminal justice context. It would have authority to hold hearings on agency AI use, and through the full Judiciary Committee, report legislation to the full House or Senate[110] establishing baseline governance requirements such as mandatory human rights impact assessments, auditability standards, and due process protections. Its work products such as formal draft

[110] *Committee Consideration*, U.S. Congress, https://www.congress.gov/legislative-process/committee-consideration.

legislation, committee reports, and white papers, could directly shape Federal procurement requirements, enforcement standards, and cross-agency guidance on AI tools used in investigations, prosecution, adjudication, and corrections.

F. Responsiveness

A dedicated Judiciary Subcommittee represents a responsive institutional mechanism for addressing the rapid deployment of AI in the criminal justice system. Because it can be created quickly under existing rules and can rely on the standing committee's staff and procedural infrastructure, it offers an immediate and actionable response to emerging harms arising from AI use across the criminal justice system. The subcommittee's capacity for frequent hearings, targeted investigations, and iterative legislative drafting allows it to adapt in real time as new technologies, failure modes, and civil rights concerns arise.

Despite strong performance on startup feasibility, transparency, policy influence, and responsiveness, this model exhibits important limitations across several criteria. First, the relevant expertise criterion is only partially satisfied: while legal and criminal justice expertise is deep, technical AI capacity is not structurally guaranteed and would depend on discretionary hiring and sustained resourcing. Second, organizational stability is inherently fragile because the subcommittee lacks statutory entrenchment, making its continuity sensitive to partisan turnover and shifting leadership priorities. These limitations suggest that while a Judiciary Subcommittee is a high-impact near-term mechanism, it may need to be complemented by more technically resourced and structurally stable entities to sustain coherent AI governance in the criminal justice system over time.

IV. Synthesis

1. Institutional Design Features

This Part draws together what the models we reviewed suggest about institutional design choices that make an AI governance entity credible, technically serious, and capable of changing practice. A few patterns show up repeatedly. Structures that are easiest to create often have the hardest time sustaining expertise, influence, and long-term institutional continuity. Bodies that produce the deepest technical work can be slow and sometimes susceptible to short-term political changes. And in many cases, the institutions that matter most in practice do not succeed because they can issue binding rules—they succeed because they develop routinized relationships with implementing agencies, maintain repeatable processes, and earn a reputation for usefulness and legitimacy. The discussion below organizes these lessons around our six criteria for success: startup feasibility, relevant expertise, transparency and accessibility, organizational stability, policy influence, and responsiveness.

A. Startup Feasibility

The main feasibility tradeoff is straightforward: the quickest-start models are often the easiest to undo. Judiciary subcommittees sit at the high-feasibility end. They can be created quickly by committee vote and can take advantage of existing jurisdiction, staff, and procedures. That makes them attractive if the goal is near-term congressional attention—hearings, agenda-setting, and legislative pressure—especially when speed matters. FACA-style advisory committees offer a similar "fast start" option. Congress, the President, or executive agencies can create them quickly, and they are a familiar way to convene diverse stakeholders without building a new agency from

scratch. The downside is that the same ease of creation can also make them easier to sunset or deprioritize.

More ambitious models require statute, sustained appropriations, and political commitment. The U.S. Sentencing Commission is a particularly strong example of a durable expert body, but its hybrid independent-agency structure—especially its placement in the Judicial Branch—was the product of legislation and would likely be difficult to replicate for AI governance as a practical matter. State sentencing commissions (including North Carolina and Minnesota) likewise require statutory authorization and continuing resources. Joint congressional committees generally require a law or concurrent resolution, which raises the barrier to entry but can also signal seriousness and build legitimacy. OTA is the cautionary example here: even an institution that is statutorily authorized and widely respected can be effectively eliminated through funding decisions. The feasibility lesson is therefore less "pick the easiest model" than "match the mechanism to the time horizon": quick-start tools for convening and agenda-setting versus institutions designed for durable, iterative governance.

B. Relevant Expertise

The most consistent predictor of success—regardless of institutional home—was whether the entity combined two forms of expertise: decision-makers who understand the domain and a standing staff that can do rigorous technical work. OTA is the clearest illustration. Its permanent staff created depth, continuity, and institutional memory that did not depend on which politicians happened to be in power. GAO shows a different tradeoff. It can be influential, and it has a science and technology assessment function (through its Science and Technology Assessment and Analytics team), but a generalist staffing model can limit depth unless outside domain-specific expertise is deliberately built into the process.

Advisory-committee models underscore both the upside and the ceiling of stakeholder diversity. The National Commission on Forensic Science assembled real cross-field expertise precisely because it brought together people from forensic science and the criminal justice system and did much of its work in public. But expertise alone did not translate into durable governance, because the body lacked permanence and depended on downstream agency uptake. The CJIS Advisory Policy Board reflects the opposite dynamic. It appears to be highly effective within its niche—state and local law enforcement information systems—likely in part because it is closely aligned with the operational needs of its implementing partners. But that alignment comes with a risk: the expertise is narrower, which may be a poor fit for AI governance questions that implicate courts, prosecution, defense, corrections, and affected communities.

Sentencing commissions show how expertise can be built intentionally through both composition and staffing. North Carolina's model uses a large, pluralistic commission to embed criminal justice perspectives directly into decision-making, paired with a smaller staff that provides research and policy capacity. Minnesota's commission is smaller and emphasizes analytic coherence and iterative adjustment, supported by an analytically oriented staff and a data-clearinghouse function—though resource limits can constrain what it can take on. The Federal commission combines a small, bipartisan governing body with a comparatively large staff and advisory infrastructure, illustrating another way to marry legitimacy and research capacity. The takeaway for AI governance is that "expertise" cannot be treated as a single variable. A workable design needs two layers: representative, credible decision-making for legitimacy and value pluralism, and standing technical capacity for evaluation, auditing, and monitoring.

C. Transparency and Accessibility

Across the models, transparency functions as a source of legitimacy. FACA is explicitly built around public accessibility, and NCFS reflected that through public meetings and dissemination of materials. Sentencing commissions likewise rely on open meetings, publicly available reports, and (in the Federal model) notice-and-comment style engagement, public hearings, and major data releases. These practices matter for AI governance because they enable external scrutiny, independent research, and meaningful stakeholder input—which are especially important in the AI context, where tools are opaque and when errors may only be discovered after harm occurs.

Legislative models also support transparency through public hearings and formal opportunities for stakeholder engagement. Judiciary subcommittees can create visible accountability quickly. OTA and GAO historically increased transparency through published reports that were usable not just inside Congress but by the public and other institutions as well. The broader lesson is that transparency is not simply a virtue; in this domain it is a functional requirement for trust and iterative improvement. At the same time, transparency that rests only on institutional culture is fragile. When transparency mechanisms are not structurally embedded, they can be narrowed or deprioritized as leadership changes or political pressures shift (as demonstrated by the Federal sentencing commission discontinuing its "open door" policy).

D. Organizational Stability

Stability determines whether an entity's work accumulates over time or resets every few years. FACA committees must be re-chartered, which can serve as an efficiency check but also makes them easier to dissolve. NCFS demonstrates the vulnerability: it built credibility and produced recommendations but lasted only four years and was phased out under a new

administration. The CJIS Advisory Policy Board offers the counterexample, having operated for more than 25 years. Its durability suggests that integration with an operational agency and a repeatable process can keep an advisory body functioning across political cycles.

Sentencing commissions bake stability into their design. The U.S. Sentencing Commission is often described as insulated from short-term political swings (based on the staggered terms of the commissioners), and the state commissions show how statutory authorization and consistent leadership can support decades of operation. Legislative models are mixed. Joint congressional committees can be durable if established by law or resolution and staffed in a bipartisan, professionalized way. Judiciary subcommittees, by contrast, are more exposed to chamber politics and can be restructured or dissolved by a new Congress. OTA remains the warning: even a respected institution can be undone through appropriations. For AI governance, stability should be treated as a design constraint, not a desirable add-on.

E. Policy Influence

Influence shows up through a variety of channels: formal authority, routinized uptake by operational actors, and reputational legitimacy. Judiciary subcommittees sit at one end of the spectrum. They can draft, amend, and report bills, giving them immediate leverage over statutory outcomes, though not necessarily over implementation. Sentencing commissions illustrate another high-impact pathway. Minnesota's guidelines are legally binding, and the Federal commission's guidelines and data infrastructure shape sentencing outcomes and broader policy debates through empirical reporting and iterative revision.

Advisory and analytic bodies can also be influential without binding authority. GAO recommendations have weight because agencies must respond and a meaningful share of recommendations are implemented, creating an accountability loop. The CJIS Advisory Policy

Board's reported adoption rate (over 96%) suggests that advisory influence can be extremely high when the body is tightly linked to an implementing agency and when recommendations are operationally feasible. By contrast, FACA committees—and NCFS in particular—illustrate the limitation of purely advisory models: recommendations require buy-in, and short time horizons and political transitions can sharply limit impact. OTA's influence was largely reputational; its reports were respected, but it lacked authority to compel action.

The practical takeaway for AI governance is that "nonbinding" does not mean "toothless." But influence depends on the pathway. Where binding rules are politically or legally difficult, influence can still be generated through required agency responses, procurement standards, conditional funding, and widely respected technical assessments. The warning is the flip side: advisory outputs without an uptake mechanism risk becoming reports that are praised but not actually implemented.

F. Responsiveness

Governing AI will require an institution that can respond to new capabilities and emergent harms. The models divide between those built for speed and those built for depth. FACA committees can be convened and re-evaluated, offering a relatively rapid way to gather stakeholder input. Legislative bodies—especially Judiciary subcommittees—can move quickly to hold hearings and draft legislation when there is political momentum, though responsiveness may be constrained by partisanship and agenda shifts.

Analytic bodies often trade speed for thoroughness. OTA's report model produced high-quality analysis but could be slow. GAO can be faster in some contexts but still faces constraints when deep expertise is required and must be sourced externally. Sentencing commissions offer a different kind of responsiveness: iterative adjustment over time. The Federal commission's annual

amendment cycle creates predictable opportunities for updates, and state commissions' regular meetings and data monitoring can detect systematic deviations and prompt revision. Responsiveness is only as strong as the underlying data systems and staff capacity.

The cross-model lesson is that responsiveness for AI governance should be designed as a portfolio: rapid mechanisms for triage and guidance when new capabilities emerge, paired with deeper evaluative processes for auditing, benchmarking, and post-deployment monitoring.

Taken together, the models reviewed here suggest that no single existing framework fully meets the demands of governing AI in criminal justice. But the design implications are fairly consistent. Durable governance requires permanence; ad hoc bodies rarely sustain monitoring and learning. It requires "two-tier" expertise—representative decision-makers for legitimacy and value sensitivity, plus standing technical and research staff for analytic continuity. It requires transparency that is structural rather than optional. It requires influence mechanisms that do not depend exclusively on formal binding authority, including required responses, conditional funding, procurement standards, and credible technical evaluation. And it requires responsiveness that is designed and resourced, with both fast and slow pathways.

Those implications push toward a hybrid approach that borrows elements from multiple models rather than copying any one of them. **In fact, the strengths of a particular criterion may often be conditional on another criterion being met—which sometimes may only be accomplished by mixing different institutional design features.** The next subpart considers vehicles for policy impact in more detail.

2. Vehicles for Impact

The institutional form of an AI governance entity matters, but so does the mechanism by which its work changes practice. Consider, for example, three vehicles for impact: (1) binding

rules, (2) nonbinding guidance paired with implementation levers, and (3) technical analyses and reporting functions that shape practice through information and reputational authority. A durable governance approach may rely on more than one vehicle simultaneously, especially given the diversity of AI use cases and the number of actors involved in criminal justice.

Binding rules can offer clarity, uniformity, and enforceability. In theory, an entity with rulemaking power could establish mandatory baseline safeguards (e.g., validation, auditing, documentation, human oversight requirements) and create consistent standards across jurisdictions. The tradeoff is feasibility and adaptability: binding rules often require higher political buy-in, can trigger legal challenges, and may be slow to update in a rapidly changing technological environment. Binding authority also raises difficult scope questions—whether the entity would regulate procurement, deployment, and use only, or also model development and training—and may be harder to apply to informal uses (e.g., individual judges or attorneys using general-purpose tools).

Nonbinding recommendations and guidance are generally more feasible and can be faster to issue and update. Guidance can take many forms: best-practice frameworks, procurement standards, model evaluation checklists, use-case-specific guardrails, and model policies for agencies and courts. The principal limitation is the risk of uneven adoption. However, nonbinding guidance can still be highly influential when paired with credible expertise, broad stakeholder engagement, and mechanisms that encourage uptake—such as required agency responses, public reporting, professional norms, or integration into procurement processes. Nonbinding guidance may also be better suited to governing "capabilities" rather than a constantly shifting set of vendor products, because it can articulate expectations for how certain functions (e.g., risk scoring, biometric identification, language generation) should be evaluated and constrained.

One implementation lever for guidance is conditional grants. A governance entity could shape practice by conditioning Federal (or state) funding on compliance with process and safeguard requirements—such as transparent procurement, disclosure of training data and limitations, pre-deployment testing for disparate impact, post-deployment auditing, logging, and correction/appeal pathways. Conditional grants are attractive because they can influence a wide range of agencies without formal regulatory control, can be tailored to different use cases, and can be updated over time. They also create incentives for agencies to build internal capacity and adopt standardized safeguards. The drawbacks are political and administrative: grant conditions can be criticized as coercive, may burden smaller jurisdictions, and require careful design to avoid encouraging superficial "check-the-box" compliance. As one of our interviewees noted, two existing models illustrate the promise of this vehicle. The Office of Community Oriented Policing Services (COPS Office) within DOJ combined grantmaking with an intellectual clearinghouse role, using funding and guidance together to shape how community policing was implemented.[111] The Law Enforcement Assistance Administration (LEAA) similarly demonstrates how large-scale grantmaking can influence local law enforcement decisions simply through the volume of resources distributed, even absent direct command-and-control authority.[112]

Finally, an entity can exert influence as a technical analysis and reporting body—an intellectual clearinghouse model. Under this approach, the entity would centralize knowledge by producing technology assessments, best-practice reports, model policies, evaluations of specific AI capabilities, and (where feasible) reproducible benchmarks and datasets. This vehicle is well suited to a fast-evolving field because it can create shared baselines for evaluation, reduce duplication across jurisdictions, and increase transparency through public-facing reporting. It can

[111] See "COPS," Community Oriented Policing Services, U.S. Dept. of Justice, https://cops.usdoj.gov/.
[112] The Law Enforcement Assistance Administration (extant 1968 to 1982) facilitated many administrative changes.

also support procurement by publishing vetted criteria and "what to ask vendors" standards. The limitation is that reporting alone may not change behavior unless tied to uptake mechanisms (professional norms, procurement requirements, funding conditions, or political accountability). In short, technical analysis is often necessary for credible governance, but rarely sufficient on its own.

In practice, the most promising strategy is mixing and matching vehicles for impact: pairing nonbinding guidance with an ongoing technical clearinghouse function and using conditional grants (or other incentives) to encourage adoption of core safeguards. Binding rules may be appropriate in narrower domains where legal authority is clear and the need for uniformity is high, but the broader governance challenge—keeping pace with evolving AI capabilities across a decentralized criminal justice system—may be best served by guidance, incentives, and credible technical evaluation working together.

3. State versus Federal Entity

A central design choice for an AI governance entity in criminal justice is the level at which it should operate: Federal, State, or some combination of both. Each option offers distinct advantages and tradeoffs with respect to legitimacy, feasibility, adaptability, and effectiveness. The models reviewed in this White Paper suggest that no single level is clearly superior across all dimensions; rather, the appropriate choice depends on the kinds of authority the entity is expected to exercise and the functions it is expected to perform.

A Federal entity offers the advantages of uniformity, scale, and coordination. Many AI capabilities used in criminal justice—facial recognition systems, risk assessment tools, language models, and analytics platforms—are developed by national or global vendors and deployed across multiple jurisdictions. A Federal body can reduce fragmentation by establishing baseline

standards, shared evaluation criteria, and common terminology. It can pool technical expertise that individual states or localities may be unable to sustain. A Federal entity may also be better positioned to engage with large technology vendors, influence national procurement practices, and set expectations that shape the broader market.

However, Federal governance also faces limitations. Criminal justice is highly decentralized in the United States, with most policing, prosecution, and sentencing occurring at the state and local levels. A purely Federal body may lack the contextual knowledge needed to shape on-the-ground practices in diverse jurisdictions. It may also face political or legal constraints that limit its ability to directly regulate state actors. In addition, a single Federal framework risks being either too general to be operationally useful or too rigid to accommodate meaningful variation in local priorities.

State-level entities offer the countervailing advantages of proximity, contextual sensitivity, and political feasibility. State commissions and agencies are embedded in the institutions that actually deploy and use AI in criminal justice. They can tailor guidance to local law and capacity; engage directly with practitioners; and build buy-in through participation by state and local stakeholders. The state sentencing commissions reviewed in this White Paper illustrate how pluralistic governance and iterative, data-driven oversight can be sustained over time within a particular jurisdiction. State-level governance can also function as a laboratory, allowing different approaches to be tested and compared.

The tradeoffs are fragmentation and unevenness. Without coordination, state-level approaches risk producing inconsistent standards and regulatory gaps. Smaller states or localities may lack the resources to sustain deep technical expertise and monitor complex AI systems. This can create inequalities in protection and oversight across jurisdictions, undermining fairness and

public trust and potentially encouraging forum-shopping by vendors or agencies seeking the least restrictive environment.

These tradeoffs suggest the value of a multi-level governance approach that combines Federal coordination with state-level implementation. Under such a model, a Federal entity could focus on setting baseline principles, developing technical standards, producing shared evaluations and benchmarks, and coordinating learning across jurisdictions. State entities, in turn, could adapt these standards to local contexts, engage practitioners and communities, and oversee implementation in light of state law and institutional realities. This layered approach mirrors successful patterns in other regulatory domains, where Federal bodies provide frameworks and resources while states retain substantial responsibility for administration and enforcement—one example is environmental law, where much of enforcement is delegated to the states.

A multi-level model also helps align institutional capacities with institutional roles. Technical assessments, vendor engagement, and cross-jurisdictional data analysis are often more efficiently handled at the Federal level, where scale and expertise can be pooled. Contextualized guidance, stakeholder engagement, and monitoring of on-the-ground effects are often better suited to state or local institutions that understand their systems and communities. Coordination mechanisms such as regular convenings would be necessary to prevent fragmentation and ensure that learning travels across levels.

In short, the choice between federal and state governance should not be framed as an either–or decision. The experience of the institutions reviewed here suggests that effective AI governance in criminal justice will likely require both centralized capacity for expertise and standard-setting and decentralized capacity for adaptation and implementation.

V. Conclusion

The central challenge of AI in criminal justice is not merely technological but institutional. Artificial intelligence capabilities are entering the system faster than existing legal structures can evaluate and govern them. The result is a growing gap between the power of these tools and the capacity of public institutions to understand their risks and ensure that they align with constitutional values, democratic accountability, and public trust.

Our analysis of existing expert bodies shows that no single institutional model is sufficient on its own. Advisory committees are easy to stand up but fragile. Analytic bodies produce deep expertise but may be slow or politically vulnerable. Legislative committees have authority but limited technical capacity. Sentencing commissions illustrate how pluralistic representation, standing research capacity, transparency, and institutional durability can be combined—but they also reveal tradeoffs in feasibility, scope, and adaptability. Taken together, these models point toward a hybrid approach.

Several design principles emerge consistently across the models. Effective governance requires sustained technical capacity alongside representative and legitimate decision-making structures. It requires transparency that is structural rather than optional, particularly where tools are opaque and stakes are high. It requires mechanisms for real-world uptake, rather than assuming that good analysis alone will change practice. And it requires institutional stability sufficient to support iterative learning, monitoring, and adaptation over time.

Finally, our analysis suggests that governance will be most effective if it is multi-level rather than monolithic. A Federal entity is well suited to set baseline standards, pool expertise, and coordinate learning across jurisdictions. State and local entities are better positioned to adapt those standards to legal context, institutional capacity, and community values, and to build legitimacy

through local participation and proximity. Designing these levels to reinforce rather than undermine each other is necessary.

However, additional questions remain for future research. For one, at what point in the AI lifecycle should governance operate? Should oversight focus only on procurement, deployment, and use, or should it also reach upstream to system design, training, and data selection? For example, should a governance body ask whether police departments should use facial recognition at all, or also examine the datasets used to train those systems and the risks they carry? Should it offer guidance on model development and validation, or confine itself to downstream safeguards? Different answers reflect different institutional ambitions and constraints, and the right scope may depend on whether a system is built in-house, purchased from a vendor, or adapted from a general-purpose tool.

A second question that future research should address is that of institutional reach. The most obvious subjects of governance are public actors: police departments, prosecutors, courts, and corrections agencies. But AI tools are also used by private actors within the criminal justice ecosystem, including defense firms, nonprofit public defender organizations, forensic laboratories, and technology vendors themselves. Should a governance entity regulate only state actors and offer nonbinding guidance to private ones? Or should certain private uses fall within its scope because of their effects on public decision making? These boundary choices shape not only the technical reach of governance, but also its legitimacy and effectiveness.

Ultimately, the core lesson is that governing AI in criminal justice is not a problem that can be solved once and for all. It is an ongoing institutional task that must evolve alongside the technologies it seeks to oversee. The goal of this White Paper is not to prescribe a single blueprint,

but to clarify the tradeoffs, illuminate workable design elements, and provide a framework for building institutions specifically suited for governing AI in criminal justice.

VI. Appendix A:

Beyond the Force of Law: On the Constitutionality of Establishing a Legislative-Branch Entity Governing the Use of Artificial Intelligence in Federal Criminal Justice:

https://law.stanford.edu/publications/beyond-the-force-of-law-on-the-constitutionality-of-establishing-a-legislative-branch-entity-governing-the-use-of-artificial-intelligence-in-federal-criminal-justice/