

CHINA'S AI REGULATIONS AND HOW THEY GET MADE

Matt Sheehan

OVER the past two years, China has rolled out some of the world's first binding national regulations on artificial intelligence (AI). These regulations target recommendation algorithms for disseminating content, synthetically generated images and video, and generative AI systems like OpenAI's ChatGPT. The rules create new requirements for how algorithms are built and deployed, as well as for what information AI developers must disclose to the government and the public. Those measures are laying the intellectual and bureaucratic groundwork for a comprehensive national AI law that China will likely release in the years ahead, a potentially momentous development for global AI governance on the scale of the European Union's pending AI Act. Together, these moves are turning China into a

laboratory for experiments in governing perhaps the most impactful technology of this era.

But international discourse on Chinese AI governance often fails to take these regulations seriously, to engage with either their content or the policymaking process. International commentary often falls into one of two traps: dismissing China's regulations as irrelevant or using them as a political prop. Analysts and policymakers in other countries often treat them as meaningless pieces of paper. President Xi Jinping and the Chinese Communist Party (CCP) have unchecked power to disregard their own rules, the argument goes, and therefore the regulations are unimportant. Other U.S. policy actors use the specter of Chinese AI governance to advance their agendas. When

Matt Sheehan is a fellow at the Carnegie Endowment for International Peace, specializing in global technology and China's artificial intelligence ecosystem. You may follow him on X @mattsheehan88. This essay is an adapted version of a July 2023 working paper, originally published as part of the "Reverse engineering Chinese AI Governance" study series of the Carnegie Endowment for International Peace.



Source: Guliver Image

The headquarters of the Cyberspace Administration of China in Beijing

Senate Majority Leader Chuck Schumer announced his plans to begin regulating AI in 2023, he described China's efforts as a "wake up call to the nation," warning that the United States could not afford to let its geopolitical adversary "write the rules of the road" for AI.

These positions are rooted in an aspect of reality, but they also create a blind spot: the regulations themselves. The specific requirements and restrictions they impose on China's AI products matter. They will reshape how the technology is built and deployed in the country, and their effects will not stop at its borders. They will ripple out

internationally as the default settings for Chinese technology exports. They will influence everything from the content controls on language models in Indonesia to the safety features of autonomous vehicles in Europe. China is the largest producer of AI research in the world, and its regulations will drive new research as companies seek out techniques to meet regulatory demands. As American- and Chinese-engineered AI systems increasingly play off one another in financial markets and international airspace, understanding the regulatory constraints and fail-safe mechanisms that shape their behavior will be critical to global stability.

And despite China's drastically different political system, policymakers in the United States and elsewhere can learn from its regulations. China's regulations create new bureaucratic and technical tools: disclosure requirements, model auditing mechanisms, and technical performance standards.

These tools can be put to different uses in different countries, ranging from authoritarian controls on speech to democratic oversight of automated decisionmaking. Charting the successes, failures, and technical feasibility of China's AI

regulations can give policymakers elsewhere a preview of what is possible and what might be pointless when it comes to governing AI.

So what do China's AI regulations contain? How did its massive party and state bureaucracies formulate them? And is it possible to predict where Chinese AI governance is headed? This paper tackles these questions using a novel approach: reverse engineering.

The approach begins with the finished product: the regulations on AI and algorithms that China has already adopted. The essay breaks down the regulations into their component parts—the terminology, concepts, and

requirements embedded in them—and then trace those components backward. It will trace their progress through China's "policy funnel" by examining the political and social roots of the ideas; how they were shaped by CCP ideology, influenced by international AI discourse, and debated by Chinese

China is the largest producer of AI research in the world, and its regulations will drive new research as companies seek out techniques to meet regulatory demands.

scholars and companies; and finally formalized by bureaucratic entities. This approach will clarify the specific aims and likely impacts of China's AI regulations and help to build a conceptual model for how China makes AI policy.

This paper provides an overview of key Chinese AI regulations to date and an introduction to the key actors and influences in the policy process.

This approach builds on the work of an international community of scholars who over the past decade have greatly improved analysis of Chinese technology policy by moving the focus further up the policy supply chain. Ten years ago, China's technology policy went largely unexamined in mainstream international discourse. Today, analysts, scholars, and the media pay much closer attention to Beijing's policy documents, often producing translations and analyses of their impact just days after their release. This approach identifies

actors from across Chinese academia, media, policy think tanks, corporations, and the party and state bureaucracies that signal and shape forthcoming AI governance. Ultimately, the aim is to both deeply understand China's existing AI regulations and to help predict what new measures may be coming around the bend.

CHINESE AI GOVERNANCE TO DATE

"AI" and "governance" are slippery concepts. Attempting to dissect all government policies that impact this basket of technologies would further muddy China's already-murky policymaking process. This essay thus focuses on a specific subset of Chinese measures: national-level policy documents that explicitly and primarily target AI or algorithms for regulation or governance.

This subset excludes several laws and regulations that impact AI development, such as the 2021 Personal Information Protection Law. It also excludes local government regulations, such as those covering autonomous vehicles, and national policy documents that focus on stimulating the AI industry rather than regulating it. The study includes some regulations that focus on algorithms rather than AI itself. It also briefly covers government some of the documents that lay out high-level guidance for the ethics and governance of AI.

Three regulations require the deep-test analysis: recommendation algorithms, "deep synthesis," and generative AI. These interconnected documents contain the most targeted and impactful regulations to date, creating concrete requirements for how algorithms and AI are built and deployed in China. Below is a brief overview of each regulation. The remainder of this essay expands on the intellectual roots and key bureaucratic actors behind these regulations.

1. Provisions on the Management of Algorithmic Recommendations in Internet Information Services

The 2021 regulation on recommendation algorithms marked the start of China's more targeted restrictions on algorithms and AI. The original motivation for the regulation was CCP concern about the role of algorithms in disseminating information online. But as that imperative worked its way through the policy community and bureaucracy, many other adjacent applications of algorithms—from setting schedules for workers to setting prices online—were tacked on. The regulation also created a reusable bureaucratic tool that would be deployed repeatedly in future regulations.

Tracing the origin of the term "algorithmic recommendation" (算法推荐) backward in Chinese state

media shows that it first emerged during a 2017 CCP backlash against ByteDance's news and media apps, in which user feeds were dictated by algorithms. The party viewed this as threatening its ability to set the agenda of public discourse and began looking for ways to rein in algorithms used for information dissemination. Much of the final regulation is dedicated to these concerns, requiring that algorithmic recommendation service providers "uphold mainstream value orientations" and "actively transmit positive energy." The regulation included some more concrete measures for online content control, such as requiring that platforms manually intervene in lists of hot topics on social media to ensure they reflect government priorities.

As policy discussions on recommendation algorithms took shape, new concerns emerged that caused authorities to add provisions addressing them. Prominent among these was public outcry over the role algorithms play in creating exploitative and dangerous work conditions for delivery workers. Similarly, as Chinese authorities cracked down on China's large tech

The technology company Tencent managed to introduce and popularize the term "deep synthesis" to describe synthetic generation of content, replacing the politically radioactive "deepfakes" with a more innocuous-sounding technical term. The new term eventually gained traction and found its way into the final regulation.

platforms during 2021, they added provisions barring providers from using algorithms for anti-competitive business practices or excessive price discrimination. Providers were also told not to build algorithms that "go against ethics and morals" by "inducing users to become addicted or spend too much." Individual users were also granted new rights by the regulation, including the right to turn off algorithmic recommendation services, to delete tags used to personalize recommendations, and to receive an explanation when an algorithm has a major impact on their interests.

Finally, the recommendation algorithm regulation created an important new tool for regulators: the algorithm registry (算法备案系统, literally "algorithm filing system"). The registry is an online database of algorithms that have "public opinion properties or . . . social mobilization capabilities." Developers of these algorithms are required to submit information on how their algorithms are trained and deployed, including which datasets the algorithm is trained on. They are also required to complete an "algorithm security

self-assessment report" (算法安全自评估报告. Here, "security," 安全, can also be translated as "safety"). Once an algorithm is successfully registered, a limited version of the filing is made public. Subsequent regulations on deep synthesis and generative AI also required developers to register their algorithms.

2. Provisions on the Administration of Deep Synthesis Internet Information Services

Around the same time as the CCP became concerned with recommendation algorithms (2017-2019), it also identified deepfakes as a major threat to its information environment and set about regulating them. During the policy incubation process, the technology company Tencent managed to introduce and popularize the term "deep synthesis" to describe synthetic generation of content, replacing the politically radioactive "deepfakes" with a more innocuous-sounding technical term. The new term eventually gained traction and found its way into the final regulation.

The deep synthesis regulation was scoped to include the use of algorithms to synthetically gener-

ate or alter content online, including voice, text, image, and video content. It requires that deep synthesis content conform to information controls, that it is labeled as synthetically generated, and that providers take steps to mitigate misuse. The regulation includes a number of vague censorship requirements, such as that deep synthesis con-

The deep synthesis regulation was years in the making, but in the end it suffered from particularly poor timing. It was finalized on November 25th, 2022, just five days before the release of ChatGPT.

tent "adhere to the correct political direction," not "disturb economic and social order," and not be used to generate fake news. When such content "might cause confusion or mislead the public," it must include a "conspicuous label in a reasonable position" to alert the pub-

lic that it was synthetically generated. The regulation also includes a number of provisions targeting misuse, such as requiring that deep synthesis users register with their real names and that platforms prompt users to obtain the consent of anyone whose personal information is being edited. Finally, it requires that deep synthesis providers make a filing to the algorithm registry.

The deep synthesis regulation was years in the making, but in the end it suffered from particularly poor timing. It was finalized on November 25th, 2022, just five days before the release of ChatGPT.

3. Measures for the Management of Generative Artificial Intelligence Services (Draft for Comment)

At first glance, China's regulatory apparatus appeared well prepared for the wave of generative AI applications that would follow ChatGPT. The deep synthesis regulation technically included most forms of generative AI, such as using the technology to create or edit images, videos, voice, and text.

But officials at the Cyberspace Administration of China (CAC) deemed the newly minted deep synthesis regulation insufficient. The core concern behind the deep synthesis measures was deepfakes, and its requirements reflect that. Requiring labels might make sense for visual or audio deepfakes, but it will not work as well for addressing new concerns around text generated by large language models (LLMs) or the increasingly general-purpose nature of the technology. In addition, the original regulation technically only covered deep synthesis services provided through the internet, leaving a regulatory gap for generative AI services that operate offline. So Chinese regulators and policy advisers quickly set to work drafting a new regulation that would

cover almost the exact same set of AI applications, but with an updated set of concerns in mind.

In April 2023, the regulators issued a draft of the new generative AI regulation for public comment. The draft reinforced many boilerplate content mandates (“embody Core Socialist

Countries and cultures may differ on the specific content of AI regulations, but they can learn from the content-agnostic structure of the regulations themselves.

Values”) and required providers to submit a filing to the existing algorithm registry. It also included several new requirements on training data and generated content that may prove extremely difficult for providers to meet.

The draft requires providers ensure the “truth, accuracy, objectivity, and diversity” of their training data, a potentially impossible standard for LLMs that are trained on massive troves of text and images scraped from millions of websites. That also poses a challenge for the draft's requirement that training data not violate intellectual property rights. The regulation mandates that generative AI not be discriminatory on the basis of race or sex and that generated content be “true and accurate,” an unsolved technical problem for LLMs that are prone to “hallucinating” inaccurate or baseless claims in their outputs.

These extremely demanding requirements for generative AI systems have

kicked off a particularly active public debate on the draft regulation. At the time of writing, Chinese scholars, companies, and policymakers are actively discussing how to maintain effective content controls without squashing China's nascent generative AI industry.

THE UNDERLYING STRUCTURE OF CHINA'S AI REGULATIONS

Countries and cultures may differ on the specific content of AI regulations, but they can learn from the content-agnostic structure of the regulations themselves. The above Chinese regulations share three structural similarities: the choice of algorithms as a point of entry; the building of regulatory tools and bureaucratic know-how; and the vertical and iterative approach that is laying the groundwork for a capstone AI law.

1. Algorithms as Point of Entry

AI governance can utilize different parts of the AI supply chain as a point of entry. Measures can focus on regulating training data, algorithms, or computing power, or they can simply impose requirements on the final

actions taken by an AI product, leaving the remedies up to the developer. China's approach to AI governance has been uniquely focused on algorithms.

China has taken a regulatory approach that is both vertical and iterative. Vertical regulations target a specific application or manifestation of a technology. This contrasts with horizontal regulations that are comprehensive umbrella laws attempting to cover all applications of a given technology.

This choice is clearly displayed in Chinese policy discourse around regulations and the decision to make algorithms the fundamental unit for transparency and disclosure via the algorithm registry. Some companies have been forced to complete over five separate filings for the same app, each covering a different algorithm used for personalized recommendation, content filtering, and more. The structure of the registry

and the required disclosures reveal a belief that effective regulation entails an understanding of, and potentially an intervention into, individual algorithms.

China's regulations are not exclusively focused on algorithms. The registry includes requirements to disclose the sources of training data, and the draft generative AI regulation has specific requirements on the data's diversity and “objectivity.” Many other requirements, such as that AI-generated content “reflect Socialist Core Values,” are defined based on outcomes rather

than technical specifics. Where regulators focus their interventions will be an important component of Chinese AI governance going forward.

2. Building Regulatory Tools and Bureaucratic Know-How

China's initial forays into governing AI have built up specific regulatory tools and broader bureaucratic know-how that can be deployed in future regulations. The algorithm registry is a standardized disclosure tool that ministries can easily include in future regulations, refining its requirements as needed. The information currently disclosed—such as data sources and security self-assessments, among others—may or may not prove to be useful to regulators. But the tool itself can act as a kind of regulatory scaffolding that eases the construction of future measures governing the technology.

Likewise, Chinese regulators are building up know-how about the technology and potential interventions. When representatives from the CAC first met with AI companies to discuss their algorithm submissions,

they reportedly “displayed little understanding of the technical details,” forcing company representatives to “rely on a mix of metaphors and simplified language.” Such meetings are an awkward but likely necessary step as bureaucrats attempt to grapple with a complex new technology. They help the regulators to build relationships to key players, to learn what they do not know, and to either upskill or hire to fill those gaps.

3. Vertical and Iterative (for now)

Stepping further back to the scope of each regulation, China has taken a regulatory approach that is both vertical and iterative. Vertical regulations target a specific application or manifestation of a technology. This contrasts with horizontal regulations, such as the European Union's AI Act, that are comprehensive umbrella laws attempting to cover all applications of a given technology. No regulation is perfectly horizontal or vertical, but most regulations lean in one direction or the other.

China's first batch of algorithm and AI regulations are relatively vertical. Each

If the government deems a regulation it has issued to be flawed or insufficient, it will simply release a new one that plugs holes or expands the scope. This iterative process can lead to confusion for companies doing compliance, but Chinese regulators view that as an acceptable cost in regulating a fast-changing technology environment.

covers a basket of related applications that Chinese regulators are concerned about and imposes requirements specific to these concerns. The baskets of applications are relatively large; for example, the recommendation algorithm regulation covers things from social media feeds to algorithms that set expected wait times for food delivery.

In addition to being vertical, the regulations are iterative. If the government deems a regulation it has issued to be flawed or insufficient, it will simply release a new one that

plugs holes or expands the scope, as it did with the generative AI draft regulation expanding on the deep synthesis measures. This iterative process can lead to confusion for companies doing compliance, but Chinese regulators view that as an acceptable cost in regulating a fast-changing technology environment.

The vertical and iterative approach of the past few years now appears to be building toward something more ambitious. In June 2023, China's State Council—the rough equivalent of the U.S. Cabinet—announced that this year it would begin preparations on a draft Artificial Intelligence Law (人工智能法) to be submitted to the National People's

Congress, China's legislature. Details remain sparse, but Chinese scholars anticipate that the law will build on the existing regulations to create a more comprehensive, horizontal piece of legislation that acts as a capstone on Chinese AI policy.

The first, overriding goal is to shape the technology so that it serves the CCP's agenda, particularly for information control and, flowing from this, political and social stability.

THE CORE MOTIVATIONS DRIVING CHINESE AI GOVERNANCE

At a high level, China's existing AI regulations are motivated by three main goals and one auxiliary goal.

The first, overriding goal is to shape the technology so that it serves the CCP's agenda, particularly for information control and, flowing from this, political and social stability. The primacy of control over information shows up clearly in the choice to first tackle AI and algorithms' influence on online content. From the CCP's perspective, for a technology to be productive it first must be tamed. As Chinese AI governance matures, this focus will likely evolve to include more industrial or security-related applications of the technology.

The second major goal behind Chinese AI governance is both obvious and frequently overlooked: to address the myriad social, ethical,

and economic impacts AI is having on people in China. The CCP prizes political control over nearly all else, but the Chinese academics, policy analysts, journalists, and technocrats who are shaping the regulations are much like their counterparts abroad—they are genuinely grappling with the diverse ways in which AI will change the lives of Chinese people. One example is in the regulatory provisions protecting workers whose schedules and salaries are set by algorithms. Chinese policy actors operate in a far more politically constrained environment than their peers in liberal democracies, with certain topics deemed taboo and many policy prescriptions off the table. But even within those constraints, there is still substantial room to explore the challenges of AI and to experiment with regulatory interventions to mitigate them.

The third goal is to create a policy environment conducive to China becoming the global leader in AI development and applications. The 2017 New Generation AI Development Plan laid out the goal of global AI leadership by 2030, which led to an explosion in industry activity and policy support for AI development. The CCP sees technology as a critical tool for boosting China's economy and national power. While the policies examined here

For the United States, one benefit of its comparatively slow progress on AI governance is the opportunity to learn from regulatory experiments abroad.

focus on regulating rather than stimulating the AI industry, the longstanding goal of AI leadership remains an important consideration shaping the regulatory debate. This is particularly prominent in the ongoing debates over how to balance the competing needs for information control and technological leadership in the draft generative AI regulation.

Finally, there remains one auxiliary goal: making China a leader in the governance and regulation of AI. American and Chinese leaders frequently point out that China has laid out some of the world's first binding regulations on AI, with the latter using it as a point of pride and the former as an impetus to action. But the rhetorical emphasis on global leadership often leads to a mistaken impression that this is a major driver of Chinese actions. An examination of the regulations and conversations with Chinese policy actors indicates otherwise. For China, being a global leader or model for AI governance is a “nice-to-have”—a small bonus for its businesses and national soft power, but not a significant driver of these AI regulations.

China's choice of first targets for regulation—recommendation algorithms and deep synthesis—indicates that

global leadership is not a core motivation for its AI governance. Recommendation algorithms are an omnipresent application of AI, but they are not a major strand of the global discourse on AI governance. If a country wanted to stake its claim to leading the world in AI governance, recommendation algorithms would not be the first target. In fact, China's regulation on recommendation algorithms does not even contain the term “artificial intelligence” in its text, despite covering many AI applications. Similarly, the term “deep synthesis” is not found in the AI governance discourse outside of China.

Chinese policy actors have even described the first-mover nature of their regulations as an added difficulty. When China began work on these regulations, the debates on the EU's AI Act were well underway in Europe, and Chinese policy analysts hoped that they could follow those debates and learn from the act. But slow progress on the AI Act meant that they had to forge ahead without the benefit of international guideposts or comparisons. For the United States, one benefit of its comparatively slow progress on AI governance is the opportunity to learn

from regulatory experiments abroad—if policymakers are willing to take foreign regulations seriously.

There is a four-layered policy funnel through which China formulates and promulgates AI governance regulations. Those four layers are real-world roots; Xi Jinping and CCP ideology; the “world of ideas”; and the party and state bureaucracies.

HOW CHINA SETS AI GOVERNANCE POLICY

This section of the essay presents a four-layered policy funnel through which China formulates and promulgates AI governance regulations. Those four layers are real-world roots; Xi Jinping and CCP ideology; the “world of ideas”; and the party and state bureaucracies. These layers are porous,

and regulations do not proceed through them in a purely linear fashion. Instead, they often pinball forward and backward through these layers, getting shaped and reshaped by academics, bureaucrats, public opinion, and CCP ideology. The order and relative importance of the layers also varies depending on the nature of the issue confronted. So far, most of the activity in the crafting of AI regulations has occurred in the third and fourth layers.

Layer 1: Real-World Roots

This layer is composed of the economic, political, social, and technological conditions that create the need for new policy and also limit

the options for regulators. Like public policy anywhere in the world, Chinese AI regulations often get their initial impetus from an exogenous shift in the real world. This can be a major evolution in technological capabilities, a new business model emerging, or a shift in underlying social or political conditions in the country. Such changes provide a spark, a problem that needs to be addressed through a change in public policy. The other components of this layer—economic, political, and social conditions—then help set the scope of what is possible with a regulation and what costs are acceptable.

In the recent draft generative AI regulation, the spark clearly derived from the leap in performance of large language models, demonstrated by ChatGPT, and the wave of public interest that followed. The policy response to that is now being shaped by factors such as China's global standing in AI and its medium-term economic growth prospects.

Layer 2: Xi Jinping and CCP Ideology

While the real-world roots provide a spark and some macro-level constraints, the second layer defines the problem and imposes its

own constraints on the policy response. In China, Xi Jinping's worldview and the CCP's evolving ideological frameworks are guides for interpreting events in the world and for deciding what constitutes a problem in need of addressing and how that problem should be understood and responded to.

Xi's decade-long and hugely successful campaign to centralize political power in his hands has led many outside observers to believe that he makes all meaningful decisions on policy and regulation.

The term "CCP ideology" is used here somewhat loosely, including not just ideology that is formally enshrined in the party's documents and ideological journals but also the broader way in which the party sees the world. The same goes for Xi and his formal

contributions to CCP ideology. He has rarely addressed specific AI regulatory issues, but the high-level priorities he sets serve as guidance for all policy actors as they address concrete issues.

This raises one of the most common misconceptions about how China sets AI policy. Xi's decade-long and hugely successful campaign to centralize political power in his hands has led many outside observers to believe that he makes all meaningful decisions on policy and regulation. Xi certainly acts as a micromanager on certain issues. Examples include giving feedback on ministry plans to crack down on the private tutoring sector, signing off

on high-level corruption detentions, and making the decision to cancel Ant Group's initial public offering after Alibaba founder Jack Ma criticized the government. Most famously, Xi tied himself directly to China's strict "dynamic Zero COVID" strategy that saw major cities locked down for months on end. When Xi takes a major interest in an issue, he can dictate policy, or at least reject versions of it that he does not like.

But it does not appear that Xi has applied this micromanagement to AI governance so far. State media have not described him as directing the regulations, as they often do in other

areas. And the regulations do not bear the normal hallmarks of an intervention by Xi: a hardline, uncompromising approach to complex policy trade-offs. Instead, provisions in the regulations can often be traced back to the work of Chinese think tanks or academics, as future papers will show.

This is not to say Xi's words do not carry tremendous power in AI policy. When he stated in a 2018 speech that China must "ensure AI is safe [or

secure], reliable, and controllable," that set up high-level goals for policymakers to strive for, while leaving the details to them. In AI governance,

Once a real-world change has thrown up an issue that needs addressing, and after the issue has been filtered through the prism of Xi Jinping and CCP ideology, it enters perhaps the most dynamic layer. This is the world of ideas, where the problem and its solution are debated by actors ranging from think tank scholars to AI scientists, and from investigative journalists to corporate lobbyists.

Xi is best thought of as setting the direction of travel for policy actors and as providing the ultimate backstop for decisions. Policymaking will broadly focus on the issues he prioritizes and take an approach resonant with his way of seeing the world. And no decision will be made that directly contradicts his expressed wishes. But when it comes to crafting Chinese AI regulations, most of the activity has so far occurred in the next two layers.

Layer 3: World of Ideas

Once a real-world change has thrown up an issue that needs addressing, and after the issue has been filtered through the prism of Xi Jinping and CCP ideology, it enters perhaps the most dynamic layer. This is the world of ideas, where the problem and its solution are debated by actors ranging from think tank scholars to AI scientists, and from investigative journalists to corporate lobbyists. This is where many policy

ideas are generated or shot down. It is where technology companies try to steer the policy dialogue in their preferred direction and where journalists can bring social issues into mainstream public discourse. While these public debates do not settle policy, they provide the intellectual grist for the bureaucratic mill.

As previously described, the debates occur within a constrained political and intellectual environment. Few of these policy actors will swim against the ideological stream, and policy solutions that contravene Xi's expressed wishes will not be entertained. How much latitude these actors have depends on the political salience of the issue at stake. For highly sensitive political issues, such as the status of Taiwan, the bounds of public discussion are extremely narrow. And what counts as political and sensitive has continuously expanded under Xi.

Nevertheless, in the area of AI regulation there is still a relatively large space for policy debates. This is perhaps due to the relatively technical nature of policies and to the freshness of the problems. How to effectively regulate AI remains a wide-open question globally, and the political interests at play in China are not yet entrenched. Ministries and state-owned enterprises have not spent decades fighting to gain leverage or to

hang onto preferential policies they have carved out. This mix of factors has made public debates over AI governance unusually lively and open.

Within that debate, several Chinese organizations and individuals stand out. Among think tanks, the China Academy for Information and Communication Technology (CAICT, 中国信息通信研究院) has emerged as particularly influential. Under the supervision of the Ministry of Industry and Information Technology (MIIT), CAICT is home to technical experts and policy analysts who have worked closely with the CAC on AI governance projects. Tsinghua University's Institute for AI International Governance (清华大学人工智能国际治理研究院) has also produced sophisticated reports drawing lessons from algorithm governance abroad and making recommendations for China. Among the many Chinese scholars contributing to the country's AI governance debates, some particularly notable individuals are Zhang Linghan (张凌寒) of the China University of Political Science and Law, Sun Ping (孙萍) of the Chinese Academy of Social Sciences, and Liang Zheng (梁正) and Xue Lan (薛澜) from Tsinghua University. Subsequent papers in this series will explore the contributions these and other scholars have made to Chinese AI governance.

Layer 4: Party and State Bureaucracies

Ideas and proposals are molded into regulations in the final layer of the policy funnel, consisting of the party and state bureaucracies. When it comes to setting AI regulation, organizations across the party and state bureaucracies are deeply interwoven. But that proximity should not be mistaken for harmonious relations. China's ministries and agencies are a notoriously "fractious and highly competitive group," always angling for their policies to be adopted at higher levels. Examining the regulations issued so far illuminates some initial conclusions about which members of this "fractious" group are prevailing in the competition for influence.

The CAC has emerged as the clear leader in the first wave of AI regulations. Tracing the roots of these regulations backward shows the CAC playing a leading role in setting the agenda and getting its pet issues in front of the highest decisionmaking bodies, such as the CCP Central Committee. When the Central Committee then approves those issues for regulation, the CAC has authored the drafts of the regulations. In writing the draft regulations, the CAC often utilizes experts affiliated with other bodies, such as scholars from think tanks affiliated with the MIIT or the Ministry of Science and Technology

(MOST). It then brings other ministries and agencies on as co-signatories when the draft becomes final, creating bureaucratic buy-in and enhancing enforcement capabilities. In this way, the CAC has acted as a hub for AI regulations.

Whether the CAC will continue to play this role remains an open question. The CAC's *raison d'être* is controlling online content, which made it a logical leader for the first batch of AI and algorithm regulations. But as AI governance shifts to other arenas such as autonomous vehicles, fintech, or frontier AI research, it is unclear if it will be able to maintain its current position in leading and coordinating the other ministries.

The MOST played a large role in early policies like the 2017 AI plan and followed that up by establishing committees and issuing high-level principles for AI ethics and governance. It also wrote the draft version of a broader technology ethics and governance measure that was later issued by the CCP Central Committee. But the MOST has taken a back seat on the more targeted regulations, not co-signing the recommendation algorithm or deep synthesis regulations. The ministry focuses primarily on issues related to research and development, making it less suited to regulating online content or certain commercial applications of AI. But the MOST's profile may rise

again as regulatory attention turns toward the underlying technology, as in the draft generative AI regulation, which imposes requirements on model training.

Beyond the CAC and the MOST, three of the more significant bureaucratic bodies are the MIIT, the Ministry of Public Security (MPS), and the State Administration of Market Reform (SAMR). The MIIT and the MPS have co-signed both the recommendation algorithm and deep synthesis regulations, while

SAMR signed only the former. Each of these organizations will likely continue playing a significant role in regulations that touch on their respective areas. The MIIT, in particular, will likely take on a greater role as AI regulation moves from online content to industrial and commercial applications of the technology.

Above all these ministries sits China's State Council and the National People's Congress. Though these organizations have not been involved in recent AI regulations, they will be the key gatekeepers for China's promised national AI law. While that legislative role confers them significant

decisionmaking power, much of the policy formulation and bureaucratic wrangling underpinning the law will likely occur within and between the subordinate ministries and administrations, particularly those listed in the preceding paragraphs.

After spending several years exploring, debating, and enacting regulations that address specific AI applications, China's policymaking community is now gearing up to draft a comprehensive national AI law.

Finally, lurking in the background are two new bodies created by party-state institutional reforms announced in March 2023: the CCP Central Science and Technology Commission (CSTC) and the National Data Administration (NDA). Neither has been formally stood up, and information on them remains scarce. The CSTC will serve as the CCP's top science and technology policymaking body. It will likely have a significant voice in AI regulatory policy, but it appears that the majority of its portfolio will focus on technology development—including major national research projects and national laboratories—rather than regulation. The CSTC will reportedly be housed in the MOST, likely giving a boost to the latter in AI governance. The NDA will focus on data infrastructure and the utilization of data to support economic and social policies. These two bodies will merit close examination as they take shape.

NATIONAL AI LAW

Chinese AI governance is approaching a turning point. After spending several years exploring, debating, and enacting regulations that address specific AI applications, China's policymaking community is now gearing up to draft a comprehensive national AI law.

That process echoes the evolution of Chinese regulations governing the internet. For much of the 2000s and early 2010s, Chinese internet governance took the form of narrow regulations issued by government ministries. As those specific internet regulations added up, the Chinese state began

formulating a wider capstone piece of legislation that would draw and build upon those regulations: China's monumental Cybersecurity Law of 2017.

China now appears to be following that same blueprint for AI, though on an accelerated timeline. There are no firm deadlines for the national AI law, but a draft version could be released in late 2023 or 2024, followed by six to eighteen months dedicated to revising the law. During that time, many of the organizations, individuals, and intellectual influences described in this paper will be shaping one of the world's most important pieces of legislation for AI governance. ●

