

Xue Lan's (薛澜) "Social Application of Artificial Intelligence" Read Within Its Own Cognitive Cage: A New Era Signification Matrix, a Two-Line Critique, and a Comparative-Regulatory Analysis

Larry Catá Backer (白栲)¹

A reflection on Xue Lan's July 3, 2026 Tsinghua keynote, read against my essays, "Modernization (现代化) and Chinese Constitutionalism's Lifeworld," "But Can You Drown a Demon 2?", and "Reflections on Measures on Cyberspace Security Supervision and Inspection by Public Security Organs"

Abstract

This essay reads a July 2026 Tsinghua keynote by Xue Lan (薛澜) — chair of China's National Expert Committee on AI Governance — as a case study in how Chinese AI-governance discourse is produced, exported, and received. It builds a matrix linking Xue's key terms to specific New Era Marxist-Leninist doctrinal referents, then pursues two lines of critique: first, an assessment, grounded in a theory of modernization as the organizing horizon of Chinese constitutionalism (drawing on my essay, "Modernization (现代化) and Chinese Constitutionalism's Lifeworld") the fundamental political line of which Xue's speech is faithful, elaborates, or deploys (finding fidelity nearly everywhere); second, a semiotic reading, using a five-stage interpretive protocol drawn from the Gerasene demoniac narrative (Mark 5), that places Xue among a corpus of AI-governance "oracles" — Palantir, Anthropic, OpenAI, Aschenbrenner, DeepSeek, and Meta — and argues his position is structurally distinct: custodial rather than supplicatory, addressed to the international order rather than a domestic sovereign. A comparative-regulatory section sets the institutional apparatus Xue's speech legitimates against the EU AI Act and the fragmented U.S. landscape. A final section, prompted by Xue's real-world engagement with Senator Bernie Sanders, identifies a deeper structure common to every framework examined — Chinese, American, and European alike — distinguishing "Left" and "Right" variants of a Leninist impulse toward concentrated, expert- or vanguard-directed control of productive and cognitive life, and argues that Xue's framework functions as a bridge between them, which is what gives the Xue-Sanders convergence its logic. The essay concludes that no framework canvassed — Chinese or Western — contemplates AI's development without some apex authority to direct it; the only contested question is who occupies that position.

Executive Summary (for the General Reader and Policymaker)

A senior Chinese AI-policy figure, Xue Lan, gave a speech in July 2026 arguing that AI's usefulness depends less on how smart the models get and more on whether societies build the surrounding infrastructure — physical, legal, and organizational — needed to put AI to work, much as the automobile only reshaped society once roads, traffic law, and gas stations were built around it. This essay treats that speech as a window into how China's AI-governance thinking works, and into what it shares with, and how it differs from, Western approaches.

¹ The AI machine system Claude was used in the preparation of this essay.

Three things stand out.

First, nearly every phrase in Xue's speech connects to a specific concept from official Chinese Communist Party doctrine — even when the speech's language sounds like neutral academic sociology of technology. This isn't necessarily insincere; it reflects how thoroughly Communist Party ideology shapes what can be said and how, in an official Chinese venue, even on a technical topic like AI infrastructure.

Second, Xue is not simply asking Beijing for permission to do things — he already speaks from within China's governing apparatus. His real audience is international: he is making a bid for China, and Chinese-style institutions, to be recognized as the model the rest of the world should adopt, rather than a rival regulatory bloc among several. This ambition showed up concretely when Xue appeared alongside American scientists at a Capitol Hill event convened by Senator Bernie Sanders in April 2026 — an event that caused real political controversy in Washington, with critics accusing Sanders of legitimizing Chinese Communist Party officials.

Third, and most broadly: this essay argues that a strain runs through Chinese AI governance, American progressive politics (Sanders' own brand of it), and Western regulatory technocracy (the EU model) alike — a shared assumption that AI's development requires some concentrated authority, whether a Communist Party apparatus, a treaty of negotiating governments and scientists, or an independent regulatory agency, to direct it. Genuinely bottom-up, undirected, or market-driven alternatives are rare in this discourse, on any side. That shared assumption, more than any disagreement over specific rules, is what allowed a Chinese Party-linked scholar and an American democratic-socialist senator to find real common ground in the same room, despite otherwise occupying opposite ends of the political spectrum.

Part I: Introduction: Context and Method

On 3 July 2026 Xue Lan (薛澜), Dean of Schwarzman College and the Institute for AI International Governance at Tsinghua, delivered remarks at Tsinghua's "2026 Academic Conference on Digital Economy Development and Governance," published on the Institute's WeChat account on July 30 and reproduced here with an English translation and framing commentary (apparently from a Substack-style newsletter, bylined Yuxuan Jia and Alex Yang). In addition to his academic appointment, Xue Lan is described as a Counsellor of the State Council of China, holds the chair of China's National Expert Committee on AI Governance and sits on the UN's Committee of Experts on Public Administration (CEPA), a UN technical advisory body that studies and makes recommendations to improve governance and public administration structures and processes for development. Prior to delivering the remarks, in May 2026, Xue joined a Capitol Hill discussion convened by Senator Bernie Sanders on frontier AI risk where, alongside other international experts, Xue Lan urged global cooperation and safety guardrails to slow down unregulated AI development.

The Remarks are worthy of critical study not merely because of the ideas they generate, but also for the context in which it was delivered and the collectives into which it was meant to be projected. This essay reads a July 2026 Tsinghua keynote by Xue Lan — chair of China's National Expert Committee on AI Governance — as a case study in how Chinese AI-governance discourse is produced, exported, and received. The remarks repay close reading not principally for its account of technology but for what it reveals about the ideological labor that Chinese academic-administrative elites are being asked to perform in the current moment of AI competition. The remarks were delivered at an official university venue, transmitted through the WeChat channel of Tsinghua's Institute for Service Economy and Digital Governance, and then translated into English with Xue's own review and revision. Every stage of that transmission chain is curatorial. One should read the text, therefore, less as an academic argument to be evaluated on its internal merits and more as a document of governance: *an artifact produced by, and for, the apparatus that is simultaneously the object and the author of its own theorization.*

Part I.A provides a summary of the remarks and my sense of their political context. Part I.B provides a note on method. Critical there is my aim to read Xue Lan's text in light of my own prior work on Marxist Leninist Lebens welt, and eventually within contemporary currents within elite circles of advancing elaboration of "left" and "right" Leninism theory within both Marxist Leninist (political vanguards) and liberal democratic (techno, expert, administrative vanguards) cognitive cages. Part II embeds Xue Lan's remarks within Chinese Marxist Leninist frameworks. Part III.A then reads Xue Lan's remarks within the conceptual universe of my paper, "'Modernization' (现代化) and Chinese Constitutionalism's Lifeworld; Constitutionalism as the Institutionalized Phenomenology of Modernization: A Semiotic-Phenomenological Analysis of Modernization Discourse in Chinese Marxist-Leninist Political-Legal Theory" [hereafter the "modernization paper"]. Part III.B then contextualizes Xue Lan's remarks within the ongoing "conversations among the leading figures of AI businesses and organs. Part IV attempts an analysis of Xue Lan's iSTS Apparatus Against Its Closest Analogues. Part V then considers the

remarks as a function of the deeper ,eta-signifier, the current global conversation about ownership and the control and direction of human activity.

Part I.A: Xue Lan's Remarks—A Preliminary Pass

Xue Lan rejects technological determinism — the view that model capability gains automatically translate into economic and social value. He argues AI has entered a "second half" in which the binding constraint is no longer algorithmic progress but the absence of a mature *sociotechnical system* around it. He borrows from STS literature (social construction of technology, large technical systems, actor-network theory, co-evolution theory) to frame technology and society as mutually constitutive rather than technology as an autonomous driver. This matters because the substantive claim of the speech — that AI's social consequences will be determined by the sociotechnical system built around it, not by model capability alone — is not merely descriptive. It is a prescription for administrative action, and the argument's rhetorical structure exists to make that prescription appear to follow ineluctably from disinterested social-scientific observation rather than from a pre-existing commitment to comprehensive state direction of technological development. This is the characteristic move of a great deal of contemporary Chinese Party-state-adjacent scholarship: the importation of Western analytic vocabulary — here, the social construction of technology, large technical systems theory, actor-network theory, co-evolutionary theory — repurposed not to interrogate power but to legitimate its exercise by a specific, named set of institutions.

Central analogy. He develops an extended automobile history in three phases: (1) 1880s–1900s, the car as an isolated technical novelty for the wealthy; (2) 1910–1945, the sociotechnical system forming around it — the Model T's price collapse, federal highway construction from 1916, gas stations, auto financing/insurance, licensing and traffic law; (3) postwar, social restructuring — suburbanization, mall/supermarket retail geography, and the oil-driven geopolitical order. His point: the car's social power came from the infrastructure and institutions built around it, not the engine alone.

Application to AI. He maps this onto three "interlocking" layers AI needs: (a) *physical infrastructure* — compute, chips (Nvidia/Intel/AMD vs. Huawei/Cambricon), hyperscale data centers, data pipelines, networking, energy; (b) *institutional infrastructure* — data circulation/ownership rules, liability regimes for AI failures (medicine, autonomous driving), auditing/certification, regulatory sandboxes, and interoperability standards across models (his "traffic signal" analogy); (c) *organizational and sociocultural systems* — workflow redesign (not bolt-on AI), human-AI task allocation that preserves human final authority, the shift from pyramidal bureaucracy to flexible team structures, and public trust/AI-literacy, citing generational and educational divides in Chinese survey data.

The automobile analogy and what it conceals. The extended automobile analogy is the speech's rhetorical centerpiece, and it is genuinely well constructed as pedagogy. Xue's three-stage periodization — isolated technical novelty (pre-1910), sociotechnical system formation (1910–1945, anchored by the Model T's price collapse and the federal highway program

launched in 1916), and social restructuring (postwar suburbanization and the petroleum-geopolitical order) — has real explanatory power, and Xue deploys it competently to make the point that computational capability without institutional and organizational complement produces underwhelming returns, a point amply supported by his own citation of AI-adoption studies showing persistent gaps between capital expenditure and realized value.

But the analogy does specific ideological work that deserves to be named. Every institution Xue lists as constitutive of the automobile's sociotechnical system — highways, licensing regimes, traffic law — was in the American case a product of federal, state, and municipal government action operating *within*, and substantially constrained by, a private capital structure that financed, manufactured, insured, and distributed the technology with only intermittent and contested state involvement, litigated through courts, legislatures, and eventually a labor movement that Xue elsewhere in the speech treats as an obstacle to efficient deployment (his aside about the truckers' union blocking autonomous long-haul trucking is offered, tellingly, without normative comment — an inconvenient friction rather than a legitimate site of contestation over who bears the costs of automation). The American case, in other words, is a story of *negotiated, adversarial, and pluralist* institution-building, not administrative design from a single directing center. Xue borrows the analogy's descriptive force while quietly substituting, for its American institutional pluralism, an implicit brief for the kind of unified, state-directed "system-building" that Chinese administrative practice already privileges. The rhetorical trick is to use a historical case whose institutions emerged from contest and bargaining to naturalize a governance model whose institutions are meant to emerge from Party-state planning.

This is worth dwelling on because it recurs as a structural feature of the argument. Xue's third layer of the sociotechnical system — the "institutional infrastructure" comprising data-circulation rules, liability regimes, auditing/certification, and "regulatory sandboxes" — is presented as if it were simply the functional equivalent of traffic law, arising to solve a coordination problem that any rational system-builder would recognize. But he then names the specific instruments: China's newly revised Cybersecurity Law, now formally incorporating AI; dedicated rules on deep synthesis and generative AI; and standards demarcating AI's "ethical and security boundaries." What is presented as sociological description is, at the level of the text's actual referents, a recitation of the current administrative-regulatory apparatus of the Cyberspace Administration of China and its sister bodies. The theoretical apparatus of co-evolution and large technical systems is doing legitimating work for a specific, already-existing, already-operating body of law — it is not a neutral analytic frame from which policy might be derived; it is a post-hoc justification for policy already made.

"Social responsibility governance" and the compliance/ethics distinction. The passage most deserving of sustained attention — and the one closest to my own long-running preoccupation with the divergence between voluntary corporate-social-responsibility regimes and hard compliance architectures — is Xue's description of what he calls the "social responsibility governance cluster" (社会责任治理集): a framework that embeds fairness, transparency, and explainability "throughout the system's lifecycle, from design requirements and runtime monitoring to institutional oversight," so that these values become, in his formulation,

"measurable and enforceable technical requirements rather than slogans that exist only at the level of principle."

Read against the background of a decade of Western AI-governance discourse dominated by voluntary principles, ethics boards, and self-assessment frameworks that have manifestly failed to constrain deployment decisions at the major frontier labs, this is a pointed — if unstated — critique. Xue is implicitly positioning the PRC's compliance-driven, statutorily grounded model of embedded technical requirements against the West's characteristic reliance on soft law, aspirational principles, and voluntary disclosure. He is right that there is a real and important distinction between principles that exist as texts and principles that exist as enforceable technical constraints instantiated in system architecture and backed by state audit power. This is a serious point, and one I have made in a different register about the global business-and-human-rights regime's chronic weakness for voluntarism.

But the comparison Xue invites elides the question that actually matters: enforceable *by whom*, against *whom*, and reviewable *through what process*. A compliance regime embedded in Party-state administrative law, where the auditing and certifying authority is itself an organ of the same political apparatus whose developmental priorities the technology is meant to serve, does not resolve the accountability problem that plagues Western voluntarism — it relocates it. The "measurability" Xue prizes is a genuine advance over toothless ethics-washing, but measurability enforced by an unaccountable regulator is not equivalent to measurability enforced by an accountable one. The speech elides this distinction entirely, treating "enforceable" as self-evidently sufficient without asking what independent check exists on the enforcer. This is not a minor omission; it is the entire question that separates a rule-of-law compliance regime from a rule-by-law administrative one, and Xue — who as chair of the National Expert Committee on AI Governance is himself embedded in the enforcement apparatus — has every institutional reason not to raise it.

"Humans retain ultimate decision-making authority": convergence of vocabulary, divergence of institution. Xue's claim that intelligent sociotechnical systems must be redesigned so that "humans retain ultimate decision-making authority and control" tracks almost verbatim the human-oversight language found throughout the EU AI Act, OECD AI principles, and most Western corporate AI-governance frameworks. This vocabulary convergence is itself a datum worth noting: it signals that Chinese AI-governance discourse, at least at the level addressed to international audiences (recall that this text was translated with Xue's active participation, evidently for export), has adopted the lexicon of the international AI-safety and AI-governance conversation almost wholesale.

But the convergence is lexical, not institutional. "Human" retaining "final authority" means something different depending on which human, situated in which institution, is empowered to exercise it. In the EU framework, however imperfectly realized in practice, the referent is meant to be the individual data subject or the operator accountable through administrative and judicial review. In Xue's framework, embedded as it is within a system whose institutional infrastructure is explicitly the Cybersecurity Law and its associated regulatory apparatus, the "human"

retaining final authority is most plausibly the organizational hierarchy — ultimately answerable to Party-state oversight bodies — rather than the individual affected by the system's operation. The speech's silence on *whose* human authority is being preserved, at a moment when it otherwise displays considerable analytic precision about layered infrastructure, is not an oversight. It is the place where the argument's universalizing vocabulary and its particular institutional home come into contact, and the text simply declines to resolve the tension by leaving the referent unspecified.

The Sanders meeting and the "public attitude" data as discourse-power moves. Two further details in the framing material deserve comment because they illuminate the speech's function rather than its content. First, Xue's participation in Senator Sanders' Capitol Hill discussion on frontier-AI risk is presented, in the accompanying commentary, essentially without gloss — as a fact about Xue's biography. But it is better read as part of a broader positioning strategy in which PRC-linked AI-governance figures cultivate visible engagement with the U.S. AI-safety-adjacent political left specifically, a constituency more receptive than either Silicon Valley accelerationism or the U.S. national-security establishment to the vocabulary of risk, precaution, and institutional constraint that Xue's speech itself deploys. This is a discourse-power move in the literal sense China's own policy vocabulary uses (话语权): the cultivation of interlocutors and venues abroad through which a PRC-aligned framing of "responsible," "sociotechnical," compliance-oriented AI governance can be naturalized as the sensible, moderate position against both unregulated market deployment and (implicitly) securitized U.S.-China technological rivalry.

Second, Xue's invocation of Chinese public-opinion survey data — showing broad public trust in AI, tempered by generational and educational stratification — performs a legitimating function for domestic audiences that should not be read naively as sociological reporting. Survey instruments administered and disseminated through state-linked research infrastructure, in a public sphere where AI-critical civil-society organizing has essentially no independent institutional base, cannot straightforwardly be read as evidence of an autonomous public disposition toward AI; they are as plausibly evidence of the success of a sustained state and platform messaging environment favorable to AI adoption. Xue uses the finding to argue for an "inclusive AI-literacy education system" that would, not incidentally, be built and administered by the same institutional apparatus responsible for the "positive attitude" being measured — a closed loop in which the state produces the disposition, measures it, and then proposes itself as the appropriate agent to manage its further cultivation.

Theory in service of system-building. None of this is to say Xue's substantive thesis is wrong. Quite the opposite, within the cognitive cages within which it was necessarily developed, and grounded in the principles of rationalizing the relationship of machine systems to humans and human collectives in which he is deeply embedded, makes sense. Moreover, the claim that AI's social and economic consequences depend on the institutional, organizational, and cultural infrastructure built around the technology, and not on model capability alone, is a defensible and increasingly conventional position in AI-governance scholarship generally, and Xue states it with

more historical texture than most Western equivalents manage. The automobile analogy, whatever its concealments, is genuinely illuminating as pedagogy.

The critical point is rather this: the speech is best understood not as an intervention in sociotechnical systems theory but as a piece of *governance theology* — an authoritative account, delivered by the state's own designated AI-governance theorist, of why the current expansion of PRC administrative control over AI development (the revised Cybersecurity Law, the deep-synthesis and generative-AI rules, the sandbox and certification apparatus) is not merely permissible but sociologically *necessary*, the inevitable institutional complement without which China's very real computational and industrial capacity in AI cannot translate into "genuine advantage." More importantly, perhaps, the remarks attempts to lay the groundwork for the extension of that argument—that the Chinese model, contextually modified, ought to serve as the global governance template for AI governance. The explicit closing move — casting the construction of this sociotechnical system as the "strategic support point for China's transition from technology follower to rule-maker" — confirms that the speech's ultimate register is not descriptive social science but competitive standard-setting: an argument, addressed as much to an international audience as to a domestic one, that the institutions China is already building are the correct institutions, and that the world's coming competition over AI will be won not at the level of the model but at the level of who successfully exports the governance architecture around it. Read this way, the text is less an analysis of sociotechnical systems than an early exhibit in the sociotechnical system it purports only to describe.

This is neither bad nor good; it is, instead, *politics*.

Part I.B: Preliminary Note on Method

Neither Xue Lan's speech nor the critique of it offered here can be read outside the ruling ideologies within which each acquires signification, interpretive power, and cognitive routing. For Xue Lan that ideology is Chinese Marxism-Leninism in its most sophisticated New Era form — a form so thoroughly institutionalized, per the modernization paper's argument, that it no longer needs to announce itself as ideology at all. It operates instead as *Lebenswelt*: the pre-theoretical horizon within which Xue's key terms already carry the century of accumulated meaning the modernization paper's nine sediments describe. Almost every key phrase in Xue's speech can be traced to a critical-word norm of the fundamental political line — not because Xue is quoting doctrine (he almost never cites it directly) but because the doctrine has become the condition of intelligibility within which a sentence about compute infrastructure or algorithmic sandboxes can mean anything at all to its intended audience.

The analysis therefore proceeds in three parts. **Part I** builds the signification matrix the framing calls for: Xue's own words and expressions, mapped directly against the key goals, initiatives, principles, and cognitive-linguistic structures of the New Era political line. **Part II** develops the two analytic lines specified: **(A)** an assessment, grounded in the modernization paper's apparatus (metasignifier, sedimentary ontology, principal contradiction, juridification), of the extent to which Xue remains faithful to, elaborates, or departs from that political line; and **(B)** a semiotic-

phenomenological reading, conducted through the Gerasene protocol developed in "But Can You Drown a Demon 2?", that places Xue among — and, as will become clear, structurally apart from — the corpus of AI oracles already assembled there. **Part III** offers a comparative-regulatory analysis, in the format of the Measures-on-Cyberspace-Security post, setting the institutional apparatus Xue's speech legitimates against its closest analogues abroad.

My "Modernization Paper" as Analytical Vehicle. Because Part III.A leans on it directly, the apparatus supplied by "'Modernization' (现代化) and Chinese Constitutionalism's Lifeworld" needs to be stated in its own terms before it is put to work.

That paper's central claim is that modernization is not one Chinese policy goal among others but the organizing horizon — the taken-for-granted *Lebenswelt*, in Husserl's sense — through which the Party and state define problems, justify institutions, and project China's future. Modernization performs this work as a **metasignifier**: not a word that empties socialism, Party leadership, reform, rule of law, common prosperity, ecological civilization, or national rejuvenation of their distinct content, but a higher-order operation of articulation that supplies the semantic field within which those otherwise heterogeneous terms acquire relational meaning and become intelligible as differentiated moments of one historically continuous project. The paper traces this horizon as a **sedimentary ontology** of nine historical layers — national vulnerability and revolutionary sovereignty, socialist construction, industrialization and the Four Modernizations, reform and opening, the socialist market economy, governance modernization, Chinese-style modernization, and national rejuvenation — that persist simultaneously rather than replacing one another, each conditioning what the later layers can mean without dictating it.

Two further pieces of that apparatus are critical to the analysis. The first is the **principal contradiction** (主要矛盾): the Party's periodic, textually documented reformulation of its own diagnosis of society's most pressing problem, which the paper identifies as the recursive mechanism by which the modernizing horizon recalibrates itself to new conditions without ever being abandoned or fundamentally questioned — most recently reformulated at the 19th Congress as the contradiction between "unbalanced and inadequate development" and the people's "ever-growing need for a better life." The second is **juridification**: the paper's six-section account of how modernization travels from Party doctrine through the Party and State Constitutions into national legislation and finally into local administrative practice, becoming at each stage not merely mentioned but constitutive of what counts as an adequate law, a capable institution, or a correct implementation. The paper's own methodological caution — that its claim is interpretive and immanent rather than empirical or causal, and that accurate reconstruction of the system on its own terms must precede any normative evaluation of it — governs the use made of the apparatus throughout Part II.A below.

Part II: The Signification Matrix

The matrix reads Xue's phrases first in their own terms, then locates the New Era referent that supplies their interpretive power, then names the doctrinal or textual source of that referent (drawing on the modernization paper's own citations where the concept overlaps), then notes the semiotic function the linkage performs.

Xue Lan's Phrase / Expression (中文 / English)	New Era Political-Line Referent	Doctrinal Anchor	Semiotic Function
人工智能发展"下半场" — AI has entered its "second half"	Periodization as a legitimating device — the New Era (新时代) itself is constituted by exactly this move: a qualitative break authorizing new leadership and new theoretical elaboration	19th Congress "New Era" periodization; modernization paper §V.A	Domesticates a technology-cycle observation into the Party's own signature historiographic gesture: history moves in stages, and naming the stage is itself an act of authority
Critique of "技术决定一切" (technology determines everything)	Refusal of vulgar/mechanical determinism; structurally parallel to the Five Features' rejection of narrowly "economistic" Western modernization	Modernization paper §V.B (Five Features, esp. Feature 3, material/spiritual coordination)	Positions AI discourse within the same "we transcend the West's one-dimensional model" rhetorical architecture used for modernization generally
"技术社会系统" / "智能社会技术系统 (iSTS)"	A domain-specific rendering of the Party's comprehensive, multi-dimensional development doctrine — the demand that no single dimension (economic, technological) stand apart from the whole	Wu Wei Yi Ti / Five-in-One overall layout (五位一体总体布局); modernization paper's account of "Chinese modernization" as sediment eight	Imports Western STS vocabulary (SCOT, LTS, ANT, co-evolution) to re-derive, from first principles, a structure the Party's own doctrine already supplies — see Part II.A
Three-layer infrastructure: 硬基础设施 / 制度与治理结构 / 组织与文化体系	Condensed technocratic restatement of the Five-in-One layout's economic-political—	Five-in-One layout; modernization paper's "juridification"	Performs the <i>translation</i> function the modernization paper assigns to juridification:

Xue Lan's Phrase / Expression (中文 / English)	New Era Political-Line Referent	Doctrinal Anchor	Semiotic Function
	cultural–social–ecological pentad, compressed to fit an engineering audience	chain (Section VI) — political theory → constitutional purpose → legislative method → administration → local practice	doctrine rendered as an operational stack a technologist can build against
算力基础设施 / "智算中心" build-out	New Quality Productive Forces (新质生产力) — Xi's 2023 formulation naming advanced technological capacity as the material substrate of the next developmental stage	Xi Jinping, Central Economic Work Conference remarks on new quality productive forces (2023); modernization paper §V.C	Places compute inside the "productive forces" slot of classical historical materialism without using the classical vocabulary — see Part II.A
数据存储中心 as "数据枢纽" / data as the core engine of AI quality	Data-element marketization reform (数据要素市场化配置) — the "Twenty Data Measures" (数据二十条) treating data as a factor of production alongside labor, capital, and land	2022 Opinions on Building a Data Basic System; broader "生产要素" (factors of production) doctrine	Extends historical-materialist factor-of-production language to a new object without naming the extension as such
"社会责任治理集" — embedding fairness/transparency/explainability as enforceable technical constraints "rather than slogans"	The recurring New Era move of converting abstract value-language into 制度化、规范化、程序化 (institutionalized, standardized, proceduralized) practice — most visible in "全过程人民民主"	Whole-process people's democracy doctrine (2021 White Paper); modernization paper's account of juridification as the process by which political	Signals that AI governance is being folded into the same legitimation architecture already built for democracy-as-procedure — implicit contrast with Western "voluntary principles"

Xue Lan's Phrase / Expression (中文 / English)	New Era Political-Line Referent	Doctrinal Anchor	Semiotic Function
	(whole-process people's democracy), which makes the identical rhetorical move for the value of democracy itself	concepts acquire enforceable legal form	governance (see first response, Part 2, §III)
Citation of the revised Cybersecurity Law, deep-synthesis/generative-AI rules, "classified, tiered management" (分类分级管理), "risk-based, agile governance"	总体国家安全观 (Holistic/Comprehensive National Security Concept) extended into cyberspace — network-power strategy (网络强国)	Cybersecurity Law (2017, rev. 2026); Holistic National Security Concept (2014 onward)	Names the actually-operating administrative apparatus (CAC and public-security organs) that the automobile analogy's "traffic law" figuratively anticipates — compare the concrete instrument examined in the Measures-on-Cyberspace-Security post
Interoperability standards / MLOps / model-service APIs as "soft-hard combined" infrastructure	National standardization strategy (国家标准化发展纲要, 2021) and the broader ambition to convert domestic technical standards into internationally adopted ones	National Standardization Strategy Outline (2021); discourse-power (话语权) doctrine	Anticipates, at the level of technical standards, the speech's closing claim about "moving from technology follower to rule-maker"
Human "retains final decision-making authority" within human-AI work systems	人民至上 (people-first / the people are supreme) — the mass-line legacy (群众路线), reformulated for an age in which the "masses" now include machine collaborators	Mass-line doctrine (Yan'an origins); contemporary "人民至上" formula in Xi-era speeches	Performs a structurally identical reassurance move to Zuckerberg's "balance of power" pledge — see Part II.B

Xue Lan's Phrase / Expression (中文 / English)	New Era Political-Line Referent	Doctrinal Anchor	Semiotic Function
Shift from pyramidal bureaucracy to project-based, team-based organization	Governance-system-and-capacity modernization (国家治理体系和治理能力现代化) — the Fourth Plenum (19th Congress, 2019) formulation	Fourth Plenum Decision (2019); modernization paper's sediment seven ("governance modernization")	Applies to private/organizational management a template the Party has already applied reflexively to itself
Positive public attitude toward AI; call for a "national AI literacy education system"	文化自信 (cultural confidence) and spiritual-civilization construction (精神文明建设) — Deng-era "two civilizations" doctrine elevated to constitutive status by the Five Features	Deng Xiaoping, "On the Reform of the System of Party and State Leadership" (1980); modernization paper §V.B, Feature 3	Recruits survey data into the same legitimating function public-opinion instruments perform throughout Party governance — evidence of the Lebenswelt's own reproduction, not a neutral finding (see first response, Part 2, §V)
Automobile narrative's shift from "catching up" (赶上) to "surpassing/pioneering" (超越)	The precise "赶超" trajectory the modernization paper identifies as the defining shift of New Era discourse — from emulation to the offer of an alternative civilizational model	Modernization paper §V.C; Belt and Road Initiative, Global Development Initiative, Global Civilization Initiative	Places AI governance inside the identical narrative arc already built for the Chinese Dream generally
Closing claim: building the sociotechnical system is China's "strategic support point for the transition from technology follower to rule-maker"	National rejuvenation (中华民族伟大复兴) — the ninth and final sediment, in which modernization becomes indistinguishable from the recovery of China's "rightful place" after	Chinese Dream doctrine (Xi, November 2012); modernization paper §V.D and its account of sediment nine	The speech's <i>telos</i> : every preceding move in the matrix converges on this single sentence, confirming that the entire address is an instance, not an exception, of the

Xue Lan's Phrase / Expression (中文 / English)	New Era Political-Line Referent	Doctrinal Anchor	Semiotic Function
iSTS's shift from "control of change" to "agility and resilience in response to change"	the "century of humiliation"		modernizing Lebenswelt
	统筹发展和安全 (coordinate development and security) — the dynamic, non-static formula that has progressively displaced an earlier, more purely defensive security posture	Holistic National Security Concept's later elaborations (post-2020); " 邢台会议 "-era formulations of dynamic risk governance	Recasts a shift in administrative philosophy (static control → dynamic coordination) as if it were a discovery about the nature of intelligent systems, when it is better read as an extension of an already-settled governance preference

Read across, the matrix does the work the framing calls for: it shows that there is scarcely a central, argument-carrying phrase in Xue's speech that does not have a ready New Era doctrinal referent standing behind it, supplying the term's actual interpretive power for the audience Xue is addressing. That density is itself the finding of Part I, and it is what makes Part II.A possible.

Part III.A: Fidelity, Elaboration, and Deviation — Reading Xue Lan Within the Modernizing Lebenswelt

The modernization paper's central claim is that modernization functions in Chinese constitutionalism not as one sign among others but as a *metasignifier*: a higher-order operation of articulation that does not replace the vocabulary of socialism, Party leadership, reform, security, or national rejuvenation but mediates among them, supplying the semantic field within which their plurality becomes an intelligible, unified project. Read against that apparatus, Xue's speech is not merely *consistent with* the New Era political line — it is a textbook performance of the metasignifier's operation, executed on a new object (artificial intelligence) that the Party's doctrinal architecture had not yet had occasion to fully absorb.

Where Xue is faithful — indeed, where he is doctrinally unremarkable. The three-layer infrastructure Xue proposes (physical, institutional, organizational-cultural) is not, on inspection, a discovery imported from Western sociotechnical-systems theory. It is a compressed restatement of the Five-in-One layout the Party has applied to every developmental object since the Hu Jintao era, now fitted to an engineering audience that would find "economic, political,

cultural, social, and ecological civilization construction" too diffuse a frame for a keynote about compute and data governance. Xue's insistence that AI's social value depends on institutional and cultural "coordination," not technological capability alone, performs — almost sentence for sentence — the Five Features' third feature: coordination of material and spiritual civilization (物质文明和精神文明相协调的现代化), the very feature the modernization paper identifies as Deng's "two civilizations" doctrine elevated to constitutive status. Xue does not cite this lineage. He does not need to. This is precisely the *Lebenswelt* effect the modernization paper describes: the interpretive rule is available to Xue's audience without citation because it has become the condition of intelligibility for any claim of this kind.

More consequential still is what Xue's speech does with the *principal contradiction* — the mechanism the modernization paper identifies as the device by which the modernizing horizon recalibrates itself without ever being abandoned. Xue never invokes 主要矛盾 by name. But the entire argumentative structure of Sections II and IV performs the identical operation at a domain-specific scale. Where the 19th Congress reformulated the national contradiction as one between "unbalanced and inadequate development" (不平衡不充分的发展) and the people's "ever-growing need for a better life," Xue diagnoses an unbalanced and inadequate relationship between AI capability, which has developed "with breathtaking speed" (日新月异), and the surrounding sociotechnical system, whose construction has been "much slower" (更加迟缓). This is the principal-contradiction operation performed in miniature, on AI as object: a diagnosis of imbalance that licenses continued Party-led system-building as the necessary corrective, without requiring any admission that the underlying developmental project — state-directed technological modernization — has failed or should be redirected. In this sense Xue is not merely faithful to the New Era line; he is *elaborating* it, supplying a template by which any future domain (biotechnology, quantum computing, whatever comes next) can be folded into the same recalibrating logic without disturbing the modernizing horizon itself.

Where the fidelity becomes more interesting — genuine theoretical elaboration. The most analytically significant feature of Xue's speech is not what it says but what vocabulary it declines to use. A speech performing this argument could have reached for the classical Marxist-Leninist categories directly at hand: productive forces (生产力) and relations of production (生产关系) would have supplied Xue's physical/institutional distinction without remainder — compute and data as productive forces, regulatory and organizational structures as relations of production that must be brought into correspondence with them lest they become a "fetter" (桎梏) on development, in the textbook Marxist formulation Mao and every subsequent CPC theorist inherited from Marx's 1859 Preface. Xue does not use this vocabulary. He reaches instead for Husserl-adjacent Western sociology of technology — social construction of technology, large technical systems theory, actor-network theory, co-evolution — terms that, read cold, could appear in an OECD or European Commission white paper without alteration.

Two readings of this choice are available, and the modernization paper's own apparatus supplies the tools to adjudicate between them. The first reading treats this as *elaboration consistent with the political line's own self-understanding*: New Era doctrine has always claimed the right to

absorb "all beneficial achievements of human civilization" (人类文明一切有益成果) as raw material for a distinctively Chinese theoretical synthesis — this is, after all, exactly what "Chinese-style modernization" as the modernization paper's eighth sediment already does with reform, opening, and the socialist market economy. On this reading, Xue's STS vocabulary is simply the newest instance of an old absorptive habit, no more remarkable than the Party's earlier absorption of market mechanisms into a still-socialist frame.

The second reading is more unsettling for the political line's own claims, and worth stating plainly: importing an academic vocabulary developed to describe technology and society as *co-constitutive*, without a governing center, risks smuggling in a structurally different account of agency than the Marxist-Leninist one it silently displaces. Actor-network theory in its home discourse is notably suspicious of privileged, controlling centers — it treats human and non-human actants as symmetrically distributed nodes in a network, precisely the kind of decentered ontology the vanguard-party model of history exists to displace. Xue's speech borrows ANT's vocabulary of mutual constitution while quietly restoring, at every turn, a controlling center (the Party-state apparatus, named concretely through the Cybersecurity Law and its associated instruments) that the borrowed theory does not itself require and, in its more radical versions, positively excludes. The modernization paper's own historical-qualification passage anticipates exactly this kind of tension: a later signifier can condense prior problematics only gradually, and the lexical, institutional, and theoretical histories of a term are not identical — their partial, imperfect convergence is itself the object worth analyzing. Xue's speech is a live instance of that imperfect convergence: an STS vocabulary of decentered co-evolution pressed into service for a recentered, Party-directed account of "coordination," with the seam between the two left entirely unexamined.

Where something closer to genuine friction — rather than mere theoretical borrowing — appears. Three moments in the speech sit less comfortably within the fundamental political line than the rest, and are worth naming precisely because Xue, as chair of the National Expert Committee on AI Governance, has every institutional incentive to smooth over exactly this kind of friction.

First, Xue's treatment of the GDPR — corrected on reflection, and on the evidence of China's own institutional discourse, from an earlier draft of this analysis that misread the move as a deviation. It is not one. China's own legislative and scholarly self-description of the *Personal Information Protection Law* (个人信息保护法, PIPL, 2021) states plainly that the statute was drafted "借鉴 GDPR 为代表的国际上其他国家或地区的有益经验" — drawing on the beneficial experience of GDPR-representative international practice — while just as consistently cataloguing where PIPL departs from GDPR once the borrowing is complete: PIPL declines to adopt GDPR's "legitimate interests of the controller" as a freestanding lawful basis for processing, recognizing only a narrower set of grounds including the specific needs of human-resources management; its consent regime is, on most comparative accounts, stricter than GDPR's; and it layers onto the European template a set of distinctly sovereign requirements — data-localization obligations for critical-information-infrastructure operators and high-volume processors, a state security-assessment pathway for cross-border transfer, and administrative

oversight lodged in the Cyberspace Administration of China — that GDPR does not require and that only make sense within China's own security-and-sovereignty framework.

That pattern is not a departure from the Five Features' "peaceful development path" polemic. It is the *precise operation* the New Era's own account of Reform and Opening (改革开放) has always required: consideration of foreign developments as a source of "beneficial experience," followed by the denaturing of the foreign element — stripped of the individual-rights-and-independent-supervisory-authority architecture that gives it its liberal-democratic character in its home jurisdiction — and its reconstitution as a sovereignty-compatible, security-embedded Chinese instrument. This is Sinification (中国化) performing its standard function on a legal transplant, not an exception to the civilizational-superiority script. Properly read, Xue's brief, admiring aside about the GDPR setting "strict external limits on data availability and use" is not evidence of comparative openness sitting athwart the political line; it is evidence of the line's own settled method for handling foreign institutional material — look abroad, extract the beneficial technique, and resovereignize it. The apparent generosity toward a Western instrument is, on this correction, further confirmation of fidelity rather than a rare instance of friction.

Second, the Sanders engagement — which requires more care than the earlier draft gave it, and which turns out, examined from Sanders' own side, to run in a different direction than a straightforward "Xue borrows Western legitimacy" reading would suggest. The April 29, 2026 Capitol Hill panel was not framed by its convener as an exercise in comparative governance appreciation. Sanders' own promotional language — "uncontrolled AI poses a severe danger to all of humanity," AI as a "runaway train," a call for "AI redlines" and eventually a "treaty to ban superintelligence" modeled explicitly on Reagan-Gorbachev nuclear arms control — situates the event inside a specifically domestic argument Sanders had already been running for months: the *Artificial Intelligence Data Center Moratorium Act*, co-sponsored with Rep. Ocasio-Cortez and framed in its own press materials as a demand that Americans "choose humanity over profit," paired with Sanders' repeated public naming of Musk, Zuckerberg, Bezos, and Thiel as the self-interested beneficiaries of unregulated AI deployment. The Chinese scientists' presence functioned, on the textual record, as a rebuttal device aimed at a specific domestic opponent: the industry-and-Treasury argument, voiced most directly by Secretary Bessent's charge that Sanders was "letting any nation other than the United States set the global standard," that any domestic AI regulation necessarily cedes ground to China. By seating Xue and Zeng Yi beside Tegmark and Krueger, Sanders sought to demonstrate that leading Chinese governance figures shared the same precautionary diagnosis — thereby converting "regulation helps China" from settled fact into contested claim. This is Sanders instrumentalizing Xue for a fight inside the American governing coalition, not Xue instrumentalizing Sanders to launder a domestic Chinese anti-monopoly critique.

The traffic was not, however, one-directional. Contemporary reporting on the transcript (Peking Nology's account of the event, corroborated by the Washington Post's critical editorial) records that Xue and Zeng used the platform specifically to argue for United Nations-centered guardrails — a multilateral rather than bilateral or U.S.-domestic framing entirely consistent with the discourse-power ambition the matrix already identifies in Xue's Tsinghua speech (rows on

standardization strategy and the "technology follower to rule-maker" closing line). Xue's own institutional purpose in accepting the invitation was therefore not to endorse Sanders' anti-oligarchy domestic program, nor to import its anti-concentration logic reflexively against China's own AI champions, but to use an American progressive senator's convening power to advance a UN-centered global-governance venue in which China's own standard-setting ambitions would have more purchase than in a U.S.-only or G7-only forum. Both men, in other words, extracted from the same room a legitimating resource for an agenda formed well before the room existed — Sanders against domestic deregulation, Xue for multilateral rule-making venues China can help shape. That the event provoked a sharp domestic backlash against Sanders (Republicans and the Treasury Secretary describing it as "partnering with foreign adversaries" and platforming "CCP officials") is itself evidence that whatever legitimacy Xue's participation conferred was neither costless nor uncontested on the American side — a datum the earlier draft's account of a smooth "legitimation move" did not capture, and which cuts against reading the episode as straightforward reputational gain for either party.

There is, however, a deeper structural reason the two men found the room congenial at all, beneath the tactical instrumentalization just described — one that Part IV develops at length: both figures, however differently positioned, are apex-direction thinkers operating within cognitive cages that treat some form of concentrated, elite-guided authority over AI's development as simply necessary, and neither is fully at home in a governance imaginary organized around genuinely diffuse, unowned, bottom-up control. That deeper affinity, rather than the tactical exchange of legitimacy documented above, is what makes the Xue-Sanders convergence more than an accident of scheduling.

Third, and most subtly: Xue's own periodization — AI's "second half," in which the constraint has shifted from technology to sociotechnical system-building — implicitly concedes that the *first half* of China's AI strategy, dominated by exactly the technological-capability race the speech now subordinates, was inadequately theorized, or at least incompletely so, precisely because the sociotechnical apparatus lagged. This is a mild but genuine act of internal criticism, performed in the idiom the modernization paper identifies as characteristic of Party self-correction: not a repudiation of the underlying project, but an admission, structurally parallel to the 19th Congress's own principal-contradiction reformulation, that a prior phase's emphasis was incomplete. Fidelity to the line, on this reading, includes fidelity to its own periodic self-critique — but the self-critique is still worth naming as such rather than absorbing silently into the speech's confident forward motion.

Summary of Line A. Xue Lan's speech is, at the level of its deep structure, an unusually clean specimen of the modernizing Lebenswelt at work: nearly every operative term in the matrix above has a doctrinal home, the principal-contradiction mechanism is performed in miniature on AI as its object, the GDPR reference turns out on inspection to be fidelity rather than friction — Sinification performing its ordinary work on a legal transplant — and the speech's closing sentence collapses the entire argument back into national rejuvenation, exactly where the modernization paper's ninth sediment predicts it must land. The genuine theoretical interest lies not in any departure from the line but in the vocabulary chosen to perform fidelity to it — an

STS lexicon whose home discourse is more decentered than the political project it has been recruited to serve. The two moments that remain genuinely worth flagging as points of friction, or at least of exposure, are narrower than the first draft of this analysis suggested: the "first half" admission, which concedes that a prior phase of Chinese AI strategy was incompletely theorized; and the Sanders engagement, which — properly read from Sanders' own promotional record rather than assumed — turns out not to be a case of Xue importing a domestic American anti-oligarchy critique into Chinese discourse at all, but a more symmetrical instrumentalization in which each side extracted a resource (Sanders: rebuttal ammunition against the "regulation helps China" charge; Xue: a UN-centered multilateral venue for China's own rule-making ambitions) for an agenda formed independently of the room, at real domestic political cost to Sanders and no clear domestic cost to Xue. That asymmetry of cost — controversy attaching to the American convener, none reported attaching to the Chinese participant — is itself a small but telling data point about which side had more to risk in being seen at the table.

Part III.B: The Semiotic-Phenomenological Reading — Xue Lan Among the Oracles

"But Can You Drown a Demon 2?" reads six contemporary voices — Palantir, Anthropic, OpenAI, Aschenbrenner, DeepSeek, and Zuckerberg/Meta — as oracles performing a five-stage Gerasene protocol (identification, naming, negotiation, authorized transfer, vessel destruction) in relation to a power each names but cannot itself contain. Xue Lan's speech belongs in this corpus. It is not a policy proposal addressed to a sovereign from outside sovereignty, as the American quintet's texts are, nor an internal testimony about a firm's operating culture, as Liang Wenfeng's remarks are. It is something the corpus has not yet had to accommodate: an oracle who speaks from within the seat of sovereign authority itself, describing the vessel that authority is already constructing. This structural fact reorganizes several of the corpus's governing categories, and is worth tracing stage by stage.

Identification. Where Zuckerberg's texts identify the animating force as an impersonal, inevitable phenomenon arriving from outside existing institutions ("superintelligence is coming"), Xue identifies AI's power as already present but *inadequately housed* — the technology has "advanced very rapidly," the problem is that "the sociotechnical systems required for its application" have not kept pace. This is a materially different identification move: not "a force is coming and must be given a body," but "a force already has bodies (models, chips, data centers) and those bodies are, for now, insufficient vessels." Xue's entire argument can be restated in Gerasene terms as the claim that the pigs currently available — mere compute, mere models — cannot yet sustain the demon at scale, and that a *larger, more elaborate herd* — physical infrastructure plus institutional infrastructure plus organizational-cultural infrastructure — must be built before the transfer can proceed safely and productively. This is architecturally the same wager Zuckerberg's revised text makes when it itemizes agents, tutors, creators, businesses, towns, watersheds, and school districts as an expanded herd of vessels. Xue simply

names the state, rather than the market and the platform, as the entity responsible for building that herd.

Naming. The corpus's central semiotic finding is that names like "superintelligence" and "open source" perform a Legion-function: appearing to designate a single referent while actually gathering multiple, at times incompatible, meanings under one lexical roof. Xue's "intelligent sociotechnical system" (iSTS) performs an identical operation, and does so quite explicitly in its own self-definition: the term is offered as distinguishing itself from ordinary sociotechnical systems precisely by naming intelligent machine agents as no longer "simple auxiliary tools" but "collaborative members of human-machine teams." This single naming act simultaneously (1) elevates the ontological status of the machine component, (2) reassures that the elevation does not threaten human "final decision-making authority," and (3) supplies, under one acronym, a bundle that ranges across semiconductor supply chains, data-center cooling systems, liability law, algorithmic auditing regimes, workplace redesign, and national scientific-literacy campaigns. Like "Legion," which answers the demand for a singular name with a name that concedes multiplicity, "iSTS" answers the demand to say what is actually happening to Chinese AI governance with a name that gathers an entire administrative apparatus under a term borrowed from academic sociology of technology. The name does the work of making heterogeneity appear coordinated.

Negotiation — or its structural absence. This is where Xue's oracle departs most sharply from every other member of the corpus, and the departure is the most important finding of this section. Aschenbrenner negotiates with the American national-security state for a "Project" not yet built. Palantir negotiates for a "technological republic" that would internalize regulatory function within itself. Zuckerberg negotiates, at extraordinary length and with increasing specificity across two texts, for a sovereign's forbearance — checkpoints offered, board review offered, but public release "not to be restricted or delayed." Even DeepSeek's Liang Wenfeng, who addresses no explicit external sovereign, performs an implicit negotiation with Party scrutiny through the vocabulary of "restraint" (克制) as strategy. Every one of these oracles speaks *to* an authority positioned, at least nominally, outside itself, requesting permission for a transfer the authority is asked to ratify rather than originate.

Xue Lan negotiates with no one. He is not asking the Chinese Party-state for permission to build compute infrastructure, liability regimes, or algorithmic sandboxes — he is describing, as an officer of the very apparatus responsible for designing them (chair of the National Expert Committee on AI Governance), a project already under construction with his own participation. The Gerasene protocol's third and fourth stages — negotiation and authorized transfer — collapse into a single administrative announcement, because the exorcist and the entity proposing where the demon should reside are, in this text, offices held by overlapping personnel. There is no equivalent, in Xue's speech, of "he gave them permission" as a distinct narrative event, because permission was never separable from the speaker's own institutional position in the first place.

This does not mean Xue's oracle performs no request at all. It means the request is displaced upward and outward, to a different addressee entirely: not the domestic sovereign, who is already possessed of (or by) the apparatus in question, but the *international order* — the community of AI-governance discourse, the audience reached through an English translation Xue personally reviewed, and, concretely, the Capitol Hill room Senator Sanders convened, where the contemporaneous transcript records Xue and Zeng Yi arguing specifically for United Nations-centered guardrails rather than for either a U.S.-China bilateral arrangement or American domestic legislation on its own terms. The request Xue's oracle makes of that audience is not "permit us to build the vessel" but "recognize the vessel we are building — and the multilateral venue we favor for validating it — as the model others should adopt": precisely the closing claim about moving from "technology follower to rule-maker," now given a concrete institutional address (the UN, rather than Washington) toward which the request actually runs. In Gerasene terms: Xue's oracle does not ask "send us among the pigs." It announces "our pigsty is already built and better engineered than yours; the exorcist's seat itself should be re-assigned — to a body neither Beijing nor Washington alone controls — according to who can demonstrate the more governable herd." This is a bid not for permission but for the office of ratifying authority, routed through a forum where China's vote counts as one among many rather than subordinate to a single rival sovereign — a considerably more ambitious oracular move than anything the American quintet or DeepSeek attempts, and one the corpus, built around a set of actors all negotiating *with* a single sovereign, had not yet had occasion to name.

Vessel destruction, and the corpus's one true believer in vessels. The Demon essay's most philosophically consequential move is its final turn: the suggestion that every oracle in the corpus, however extreme its disagreements with the others, shares one unexamined premise — that the demon must have an address, a body, a vessel, and that governance is therefore always a question of who controls that vessel rather than whether the entity requiring control can be addressed as located at all. Read against that closing provocation, Xue Lan is not merely one more oracle sharing the corpus's blind spot. He is the corpus's *purest* case of vessel-commitment, because his entire theoretical and political project depends on the proposition that AI's power is *nothing apart from* its vessel. Xue's central empirical claim — that most AI deployments produce underwhelming returns because "只听楼梯响，不见人下来" (all the noise of footsteps on the stairs, no one actually comes down) — is precisely the claim that capability without adequate housing is capability without effect. Where the Demon essay's closing section asks what happens to accountability once intelligence no longer needs a body, Xue's entire national strategy is staked on the opposite wager: that intelligence's *entire social significance* is a function of the adequacy of the body built for it, and that the Party-state's comparative advantage lies precisely in its capacity to build better, more comprehensive, more integrated bodies — physical, institutional, and cultural at once — than a market-fragmented alternative can.

This makes Xue's oracle, among the seven now assembled across the two readings, the one for whom the Demon essay's unvesselized prospect would be most devastating if true. Every other oracle loses an accountability relationship if the vessel dissolves; Xue loses the entire theoretical ground of Chinese-style modernization's claim to AI leadership, because that claim rests on the proposition that comprehensive, state-coordinated vessel-construction — not raw model

capability — is where genuine competitive advantage will be decided. If intelligence turns out not to need a body at all, the sociotechnical-systems thesis does not merely require revision; it loses its object entirely, and with it the "strategic support point" the speech's closing sentence stakes an entire national project on. The automobile analogy itself quietly concedes this vulnerability: the essay's own logic (Part 3, §II above) already showed how the analogy naturalizes a state-directed model by drawing on a case of institutional pluralism. Read now through the Gerasene lens, the choice of the automobile — the paradigmatic *bodied, located, infrastructure-dependent* technology — as the master analogy for AI is itself a further symptom of the same commitment: Xue reaches, of every possible historical parallel available to him, for the one technology whose power is most completely inseparable from roads, stations, and license bureaus. A demon that needs no vessel would falsify the analogy at its root.

Part IV: Comparative-Regulatory Analysis — Xue Lan's iSTS Apparatus Against Its Closest Analogues

Following the format of the Measures-on-Cyberspace-Security comparison, this section treats Xue's "institutional infrastructure" layer — not as abstract theory but as shorthand for an actually-operating and still-forming regulatory apparatus (the revised Cybersecurity Law's AI provisions, dedicated deep-synthesis and generative-AI rules, "classified, tiered" risk management, regulatory sandboxes, and the standards regime governing model interoperability) — and sets it against the closest functional analogues in the United States and European Union. As in that earlier comparison, statements about relative "breadth," "coherence," or "proactivity" are interpretive judgments, not objective rankings; each system encodes different constitutional and political commitments the comparison cannot itself adjudicate.

China's apparatus, as Xue's speech describes and legitimates it, gives the Party-state, acting through the Cyberspace Administration of China, the public security organs, and sector regulators, integrated authority over AI's physical layer (compute, data, energy allocation, largely through industrial policy and state investment), its institutional layer (the Cybersecurity Law's AI provisions, deep-synthesis and generative-AI-specific rules, standards defining "ethical and security boundaries," classified/tiered risk management, and regulatory sandboxes), and — distinctively, and without a close analogue elsewhere in this comparison — its organizational-cultural layer, through state-sponsored AI-literacy campaigns and an explicit doctrine that workplace and organizational redesign around AI is itself a governance object, not merely a private managerial choice.

Dimension	China (per Xue Lan / PRC apparatus)	United States	European Union
Locus of authority	Centralized, Party-led: Central Cyberspace Affairs Commission	Fragmented: no single federal AI statute; sectoral regulators (FTC, FDA,	Centralized at the regulatory-instrument level

Dimension	China (per Xue Lan / PRC apparatus)	United States	European Union
Physical/compute infrastructure policy	overall direction; CAC content/algorithm authority; public security organs' inspection and enforcement power (cf. MPS Order No. 176's extension of this model to cyberspace security generally); sector regulators subordinate to unified coordination	financial regulators), a patchwork of state AI laws, and executive-branch policy instruments (shifting across administrations) operating without a unifying administrative center	but civilian and multi-institutional: EU AI Act enforced through national competent authorities and the new AI Office, alongside GDPR supervisory authorities and (for infrastructure-adjacent matters) NIS2 competent authorities
	Explicit state-directed industrial policy: intelligent-computing-center build-out, chip self-sufficiency drive (Huawei, Cambricon named directly against Nvidia/Intel/AMD), energy allocation increasingly treated as a strategic AI-adjacent resource	Market-led with targeted federal intervention (CHIPS Act incentives, export controls on advanced chips to China); no unified national compute-allocation authority	Market-led with EU-level coordination ambitions (European chip sovereignty initiatives, the AI Continent Action Plan and successor efforts) but implementation remains substantially a Member-State competence
	Revised Cybersecurity Law formally incorporating AI; dedicated deep-synthesis and generative-AI rules; standards defining ethical/security boundaries; "classified, tiered" (分类分级) risk	No comprehensive federal AI statute; risk addressed through existing sectoral law (FDA for medical AI, financial regulators for algorithmic trading, state tort law for autonomous-vehicle liability), voluntary frameworks (NIST AI RMF), and a	EU AI Act: comprehensive, risk-tiered (prohibited / high-risk / limited-risk / minimal-risk), horizontally applicable across sectors, with conformity-assessment, post-
Institutional/regulatory infrastructure			

Dimension	China (per Xue Lan / PRC apparatus)	United States	European Union
Organizational/cultural infrastructure as an explicit governance object	management; regulatory sandboxes; liability regimes still developing, particularly for autonomous systems and medical AI	shifting executive-order landscape	market monitoring, and significant fines; GDPR continues to govern the data layer
	Distinctive and largely without analogue: state-sponsored AI-literacy campaigns tied to survey instruments (14th Survey on Scientific Literacy), explicit doctrine that workplace redesign around human-AI collaboration is a matter for coordinated national attention, not solely private management discretion	Absent as a distinct regulatory category; workplace AI adoption treated as a private managerial and, at most, labor-law matter (NLRB guidance, state legislation on automated employment decisions)	Emerging but narrower: AI Act's transparency and human-oversight requirements touch workplace deployment (particularly for high-risk employment-related systems), but there is no equivalent of a state-sponsored national AI-literacy campaign tied to public-opinion measurement
	"Humans retain ultimate decision-making authority and control" — textually convergent with Western human-oversight language (see first response, Part 2, §IV) but institutionally embedded in a system where the "human" exercising authority is, in the first instance, organizational and	Sectoral and inconsistent: some domains (medical devices, credit decisions) impose explicit human-review requirements; no general federal human-oversight mandate	AI Act Article 14: explicit, horizontally applicable human-oversight requirement for high-risk systems, with named mechanisms (ability to override, ability to intervene, ability to decide not to use the system)
Human-oversight formula			

Dimension	China (per Xue Lan / PRC apparatus)	United States	European Union
	state-linked rather than the individual data subject		
Proactivity / temporal horizon	Continuous, coordinated, and explicitly staged: Xue's own "second half" periodization signals a deliberate, Party-announced shift in developmental emphasis, consistent with the two-centenary-goal and 2035/2049 staging the modernization paper documents	Largely reactive at the federal level, punctuated by executive-branch initiatives of uncertain durability across administrations; more continuous at state level (California, Colorado, others) but uncoordinated nationally	Continuous and staged by design: AI Act's phased entry into force (prohibited practices first, high-risk obligations later) mirrors, at the regulatory-instrument level, a similar staged-implementation logic to China's, though without China's civilizational periodization rhetoric
Legitimacy function	National rejuvenation and discourse-power (话语权): the apparatus is legitimated not only by domestic welfare claims but by the explicit ambition to become the model other states adopt — see the speech's closing line and Part II.A above	Domestic: innovation leadership, national-security competitiveness (particularly vis-à-vis China), consumer protection; little explicit ambition to export the regulatory model itself as a civilizational offering	Explicit ambition to export the AI Act as a global standard (the "Brussels effect"), legitimated through rights-protection and democratic-accountability language rather than civilizational-rejuvenation language
Oracular/vessel structure (per Part II.B)	Custodial: the state apparatus and the entity narrating it substantially overlap; no separate sovereign	Supplicatory at the level of private oracles (Zuckerberg, OpenAI, Anthropic all negotiate with a sovereign	Regulatory-custodial but civilian: the EU institutions occupy the exorcist's seat

Dimension	China (per Xue Lan / PRC apparatus)	United States	European Union
Extraterritorial ambition	from whom permission is sought; the "request," where one exists, runs outward to the international order for recognition as rule-maker	positioned as external); the state itself, however, occupies a comparatively reactive rather than custodial position relative to those private oracles	directly, but through independent agencies rather than a Party apparatus whose personnel (like Xue) simultaneously narrate and constitute the system
	Explicit and civilizational: Belt and Road, Global Development Initiative, and Global Civilization Initiative extend the "alternative path" offer to the Global South; the AI sociotechnical-systems thesis is offered as exportable theory, not merely exportable product	Primarily commercial/competitive (export controls, chip restrictions aimed at containing rivals) rather than institutional-model export	Explicit and regulatory: the AI Act is designed, and has already functioned in the GDPR precedent, to extend EU standards extraterritorially through market access conditions

Framing caveats, in keeping with the format of the Measures comparison:

- Xue's speech describes an apparatus still substantially in formation; the deep-synthesis and generative-AI rules it cites, like the cyberspace-security inspection measures examined separately, are recent instruments whose enforcement practice will only become legible over time. Claims here about "coordination" and "reduced burden" repeat the apparatus's own self-description more than they confirm its operation in practice.
- The organizational-cultural layer is the dimension in which China's approach is least commensurable with either the U.S. or EU comparators — not because those systems lack workplace AI concerns, but because neither treats national AI-literacy cultivation as a coordinated state project tied to public-opinion measurement and legitimation. Whether this represents governance innovation or ideological reproduction (see first response, Part 2, §V, and Part II.A above) is a normative question this comparison does not resolve.
- The EU AI Act is the closest functional analogue to China's institutional layer in structural terms (comprehensive, risk-tiered, horizontally applicable) but rests on an

entirely different legitimacy architecture — rights-protection and civilian regulatory independence rather than civilizational rejuvenation and Party-state coordination. Structural similarity at the level of regulatory technique should not be read as convergence at the level of political function.

- The United States remains, as in the cyberspace-security comparison, the most fragmented of the three systems, and the comparison's greatest source of uncertainty: federal AI policy has shifted meaningfully across recent administrations in ways state-level and sectoral instruments have not, making any snapshot comparison more provisional here than for either China or the EU.
- Where Xue's apparatus and the EU's regulatory ambition meet — both are self-consciously exportable models seeking adoption beyond their home jurisdiction — the resulting friction is unlikely to remain confined to technical standard-setting bodies. Both the modernization paper's account of discourse power and the Demon essay's account of oracular permission-seeking suggest that the coming contest will be fought as much over who occupies the position of *rule-maker* — who gets to narrate the exorcist's role rather than merely request it — as over the content of any particular rule.

Part V: The Deeper Meta-Signifier — Apex Ownership, Left and Right Leninism, and the Logic of the Xue-Sanders Convergence

Pulling the preceding analysis together — Xue's own approach, the critique of it, and the critique's own necessary engagement with the wider corpus of Western AI-governance proposals in Part III.B — exposes a structure more fundamental than the modernization metasignifier that organized Part III.A, and one that runs beneath nearly everything said so far without having yet been named directly. Stripped of its considerable local complexity, the structure concerns a single question: *who or what is entitled to own and direct the productive forces of a society — understood broadly enough to include not only capital and technology but ideas, culture, values, behavior, and the very shape of collective aspiration — and toward what end, even an end as undefined as "collective self-fulfillment"?*

The genealogy of the question. In the era in which Rosa Luxemburg wrote — and as a kind of homage to the limits already visible in Marx's own text — this question was answered almost entirely in the vocabulary of industrial production, on the assumption that everything else (culture, ideals, values) would follow from, or was reducible to, class consciousness understood as a function of one's relationship to the means of production. The resulting language was anti-"capitalist" in its surface form — as it remains today among certain factions of Western intellectuals and the propaganda organs of Marxist-Leninist states — but at its structural core it shared with its liberal opponents a common premise: that capital, as object and as productive force, properly occupies a sovereign, organizing role in social life. The dispute was never really over whether such an apex-organizing role should exist. It was over who should occupy it. The

liberal-capitalist tradition answered with an ideal of diffuse ownership, aiming eventually at the utopia of everyone owning "it" — shares, property, capital broadly distributed. The Marxist tradition answered with the State as sole capitalist, the singular owner acting in trust for "the workers," undertaking the leading role in moving the collective it embodied toward communism.

The totalizing expansion. The Soviets, and the Chinese Communist Party considerably more so, expanded this apex-ownership logic well beyond production narrowly conceived, to encompass every aspect of human relations, conditions, cognizance, and behavior. This is precisely the totalizing move the modernization paper's sedimentary ontology traces domestically, and that Part II.A of this analysis found performed, in miniature, on AI as a new object of Party-directed system-building: not merely the ownership of factories and compute clusters, but the direction of organizational culture, social trust, AI literacy, and "correct" human-machine collaboration — the entire "sociotechnical system" Xue proposes is, read through this lens, a bid to extend apex direction from the means of production to the means of cognition itself.

The Western mirror: two Leninisms. What is easy to miss, because the vocabulary differs so completely, is that Western intellectual and political life has produced its own versions of this same apex-ownership impulse, extending the claim of legitimate direction beyond production into culture, meaning, and behavior — without needing a Leninist party-state to do it. Two variants are worth distinguishing, however imperfectly they map onto the American and European political spectrum.

"Left" Leninism is by now a reasonably well-defined phenomenon: a vanguard sensibility, lodged not in a party-state but in universities, media institutions, professional and human-resources bureaucracies, and philanthropic and non-governmental structures, that claims the authority to direct consciousness, language, and behavior toward a substantively egalitarian or liberationist end — without owning literal productive capital, but exercising something functionally continuous with the vanguard party's claim to know, and to enforce, the correct direction of social development. It shares with classical Leninism the premise that a self-selected, credentialed elite is positioned to discern and administer the collective's proper trajectory; it differs chiefly in dispensing with the party-state apparatus and substituting institutional, cultural, and reputational enforcement mechanisms for coercive state ones.

"Right" Leninism is more diffuse and, tellingly, largely unable to name itself — because the postwar collapse and disgrace of the century's actual fascist movements evacuated the available vocabulary for any avowedly elite-directed, anti-liberal politics of the right, leaving "everything fascist" as a term of pure abuse rather than a describable political position. What survives as an articulable non-fascist right crystallizes, in its purest self-description, as free-market ideology and bottom-up, iterative development — a rejection of "modernization" in the constructivist, top-down sense, in favor of spontaneous, emergent order. But this self-description covers only part of the actual terrain. Right Leninism's more consequential manifestations are technocratic rather than libertarian: expert-guided administrative governance (a "right" vanguard of credentialed specialists rather than a Party cadre), the triumph of scientism as a mode of legitimation, and the European tradition of independent regulatory supervision — all of them aimed, like their left

counterpart, at a kind of perfectionism, but shorn of Communist determinism and redirected toward "civilization," institutional stability, or some other collective good left conveniently unspecified. Modern administrative bureaucracies, public and private, are among the principal sites where this contest is worked out in practice; supranational bodies such as the European Union, and to a lesser extent the United Nations system, are unusually sophisticated sites for the elaboration of Right Leninism grounded in a liberal-democratic vanguardism — expert bodies claiming the authority to define, monitor, and enforce the "correct," risk-managed, rights-respecting trajectory of technological and social development on behalf of publics who have delegated, rather than directly exercised, that authority.

Where Chinese Marxism-Leninism actually sits. Understood this way, Chinese ML aligns conceptually far more closely with Right Leninism than with free-market liberal democracy, which Chinese doctrine consistently characterizes — and genuinely perceives — as a source of chaos and instability, the two conditions the New Era political line treats as the paramount dangers to be foreclosed (recall the "coordinate development and security" formula and the principal-contradiction mechanism discussed in Part III.A). Both Chinese ML and Right Leninism share a commitment to expert, apex-directed guidance of technological and social development toward a stipulated collective good, mediated through credentialed administrative authority rather than market spontaneity or unmediated democratic contest; where they part company is in Chinese ML's explicit Party-state ownership of the apex position and its residual Marxist teleology (however attenuated), against Right Leninism's diffusion of that apex position across nominally independent regulatory and technocratic bodies answerable, in principle if not always in practice, to liberal-democratic electorates.

Xue Lan as bridge. Xue's speech, and his engagement with Sanders, can now be read as a deliberate — or at least structurally intelligible — bridging move between Left and Right Leninism in the specific domain of AI governance. His "sociotechnical systems" framework appeals simultaneously to the vocabulary of liberal-democratic vanguardism (citing the GDPR approvingly, gesturing toward international standard-setting bodies, favoring UN-centered multilateral guardrails over bilateral or unilateral American control, using the "risk-based, agile governance" and "regulatory sandbox" vocabulary that is the house style of Brussels-style Right-Leninist technocracy) and to the techno-bureaucratic guidance tradition of the Chinese mass line, contextually updated for AI (human-machine "teams" that nonetheless preserve human, meaning institutional, "final decision-making authority"; state-sponsored AI-literacy campaigns; organizational redesign as a matter for coordinated national, not merely private, attention). Xue's framework is legible, and even attractive, to a Right-Leninist Brussels regulator and to a Party cadre alike, because it is built to be legible to both — it is, in this specific and narrower sense, doing at the level of AI-governance discourse exactly what the modernization metasignifier does at the level of Chinese domestic doctrine: mediating among otherwise heterogeneous vocabularies without requiring any of them to abandon their own internal logic.

This is also what makes the Sanders convergence more than an accident of scheduling or a pair of tactically compatible press events, as Part III.A's narrower account already suggested but did not fully explain. Sanders' own politics is itself a fusion point of Left and Right Leninist

impulses on the American side: his rhetoric of anti-oligarchy — the naming of Musk, Zuckerberg, Bezos, and Thiel, the demand to "choose humanity over profit" — is recognizably Left-Leninist in its moral vocabulary of illegitimate concentrated ownership and its appeal to a mobilized collective interest against capital. But his preferred remedy is not diffuse, bottom-up redistribution or worker ownership; it is expert-mediated, state-to-state elite control — "AI redlines" negotiated by "top scientists," a "treaty to ban superintelligence" modeled explicitly on Reagan-Gorbachev arms control, a moratorium administered by federal regulators pending the construction of "national safeguards." That remedy is recognizably Right-Leninist in its structure: an apex of credentialed experts and negotiating states, not a market and not a mobilized public, is to hold and direct the dangerous capability. Sanders' politics, in other words, already fuses the two Leninisms this section has distinguished, which is precisely why a Chinese interlocutor whose own framework performs an analogous fusion — mass-line rhetoric married to techno-bureaucratic, internationally-legible governance vocabulary — found the room so congenial, and why the alignment, ephemeral and tactically fraught as the surrounding controversy shows it to be, was not merely coincidental.

Why this matters for the analysis as a whole. This deeper meta-signifier does not replace the modernization-metesignifier reading of Part II.A; it explains why that reading was necessary but not sufficient. Modernization-as-metesignifier accounts for Xue's fidelity to Chinese domestic doctrine. The apex-ownership meta-signifier accounts for why his framework travels so easily beyond China's borders and finds real purchase among audiences who share none of that domestic doctrine's Marxist commitments but share fully its underlying premise that AI's development requires direction by some concentrated, credentialed authority rather than by diffuse, unowned, spontaneous processes. It also supplies the missing connective tissue between Part III.A and Part III.B: the Gerasene corpus's shared, unexamined assumption that the demon must have a vessel — some locatable body it can be housed in and, in principle, governed through — is the semiotic-phenomenological face of the same structure examined here in its political-economic form. Every oracle in that corpus, Zuckerberg's nominally market-diffuse platform included, ultimately proposes an apex — a board, a Project, a Party, a firm's internal vision, a treaty regime of negotiating states — that will hold and direct the force in question. None of the governance frameworks canvassed anywhere in this analysis, Chinese, American, or European, genuinely contemplates AI's development proceeding without some form of concentrated, apex-level ownership or direction; the dispute, in every case examined, is only ever over who or what should occupy that position, never over whether the position should exist at all. Xue Lan's bridging appeal to both Left and Right Leninist audiences, and its natural resonance with a senator whose own politics already fuses the two, is therefore not an anomaly within the corpus but its most legible symptom.

VI: Closing Synthesis

Read together, the four parts of this analysis converge on a single finding, approached from four directions. The matrix shows that Xue's speech is saturated with New Era doctrinal content beneath an STS-inflected surface vocabulary. The two-line critique shows, first, that the saturation is faithful almost to the point of orthodoxy, elaborated chiefly through a choice of

borrowed vocabulary whose home discourse sits in mild tension with the recentered political project it has been recruited to serve; and second, that Xue's oracular position is structurally unlike every other voice in the comparative corpus, because he speaks from inside, rather than in supplication to, the sovereign authority the other oracles are all still negotiating with — his true request is addressed not downward to a domestic regulator but outward to the international order, for recognition as its rule-maker. The comparative-regulatory analysis shows that this same custodial, civilizational ambition is not confined to rhetoric: it is already instantiated in an operating apparatus whose closest functional analogue, the EU's AI Act, shares its structural comprehensiveness but not its legitimating grammar. And Part V shows that beneath even these three findings lies a still deeper structure common to every framework this analysis has examined — Chinese, American, and European alike — in which the only live question is who or what should occupy the apex position directing AI's development, never whether such a position should exist at all.

What remains most exposed, across all four readings, is the wager the speech stakes everything on: that AI's power is entirely a function of the adequacy of the body built to house it, and that the Party-state's capacity to build a more comprehensive body than any market-fragmented rival constitutes China's decisive advantage. It is a wager made with unusual theoretical sophistication, and one that finds unexpected company — in a Brussels regulator's risk-tiered technocracy, in a Left-Leninist senator's call for expert-negotiated redlines, in a market oracle's founder-controlled safety board — among interlocutors who share almost nothing else with Xue Lan except the conviction that some apex must hold the reins. It is also, on the terms of the corpus this analysis has now placed Xue within, the wager most vulnerable to the one question none of the seven oracles — Chinese, American, Left-Leninist, or Right-Leninist — has yet allowed itself to ask.

Discussion Draft prepared for further development; citations follow the conventions of the source materials referenced above.

References

Backer, Larry Catá. "But Can You Drown a Demon-2?: Reflections on Mark Zuckerberg — 'The Future Is for Everyone' (10 August and Its July 28 Predecessor) and the Oracular Discourses of the AI Vanguard." Working Paper Essay, Penn State Dickinson Law, August 10, 2026. https://www.backerinlaw.com/Site/wp-content/uploads/2026/08/FINAL_But_Can_You_Drown_a_Demon-2.pdf.

Backer, Larry Catá. "'Modernization' (现代化) and Chinese Constitutionalism's Lifeworld; Constitutionalism as the Institutionalized Phenomenology of Modernization: A Semiotic-

Phenomenological Analysis of Modernization Discourse in Chinese Marxist-Leninist Political-Legal Theory." Working Paper Draft, Penn State Dickinson Law, August 2026.

Backer, Larry Catá. "Reflections on Measures on Cyberspace Security Supervision and Inspection by Public Security Organs (《公安机关网络空间安全监督检查办法》); Text and Comparison with U.S., E.U., and Brazilian Measures." *Law at the End of the Day* (blog), August 16, 2026. <https://lbackerblog.blogspot.com/2026/08/measures-on-cyberspace-security.html>.

Backer, Larry Catá. "The Semiotics of a Systems Theoretical Application of Chinese New Era Marxist-Leninism in AI Systems for the World — Reflections on 习近平在 2026 世界人工智能大会暨人工智能全球治理高级别会议开幕式上的主旨讲话 [Keynote Speech by Chinese President Xi Jinping at the Opening Ceremony of the 2026 World AI Conference]." *Law at the End of the Day* (blog), July 21, 2026. <https://lbackerblog.blogspot.com/2026/07/the-semiotics-of-systems-theoretical.html>.

Backer, Larry Catá. "Reflections on 习近平：在庆祝中国共产党成立 105 周年大会上的讲话 [Xi Jinping: Speech at the Gathering Celebrating the 105th Anniversary of the Founding of the Communist Party of China]." *Law at the End of the Day* (blog), July 2026. <https://lbackerblog.blogspot.com/2026/07/reflections-on-105-xi-jinping-speech-at.html>.

Backer, Larry Catá. "Promoting High-Quality Development of..." *Law at the End of the Day* (blog), June 2026. <https://lbackerblog.blogspot.com/2026/06/promoting-high-quality-development-of.html>.

Backer, Larry Catá. "Reflections on 张冠梓: 从世界历史纵深把握中国式现代化的时代价值 [Zhang Guanzi, Grasping the Contemporary Value of Chinese Modernization from the Depth of World History]." *Law at the End of the Day* (blog), May 2026. <https://lbackerblog.blogspot.com/2026/05/reflections-on-zhang-guanzi-grasping.html>.

Bessent, Scott. Statement on X regarding Senator Bernie Sanders's Capitol Hill AI forum. Quoted in Chris Katje, "Scott Bessent Blasts Bernie Sanders for Hosting Chinese Scientists for Discussions on AI Threats," *Benzinga*, April 28, 2026. <https://www.benzinga.com/markets/tech/26/04/52083312/scott-bessent-blasts-bernie-sanders-chinese-ai-scientists-hugo-chavez-economy-criticism>.

Chan, Kelly. "U.S., Chinese Scientists Rebut Narrative on Regulation and the 'AI Race' at Sanders Event." *InsideDefense.com*, May 1, 2026. <https://insidedefense.com/ai-news/us-chinese-scientists-rebut-narrative-regulation-and-%E2%80%98ai-race%E2%80%99-sanders-event>.

"Bernie Sanders Faces Criticism over Capitol Hill AI Panel Guest List." *Harian Basis*, April 28, 2026. <https://www.harianbasis.co/en/sanders-criticized-chinese-ai-panel>.

"Bernie Sanders Set to Speak alongside CCP Officials at Capitol Hill AI Panel." *Fox News*, April 27, 2026. <https://www.foxnews.com/politics/bernie-sanders-plans-schmooze-with-top-beijing-ai-experts-ignites-backlash>.

"Bernie Sanders' Plans to Schmooze with Top Beijing AI Experts Ignite Backlash: 'Holy S---'." *AOL*, April 27, 2026. <https://www.aol.com/articles/bernie-sanders-plans-schmooze-top-220045829.html>.

China Law Insight. "个人信息保护立法效果、理念及价值平衡——欧盟 GDPR 生效实施三周年比较与前瞻" [Effects, Concepts, and Value Balance of Personal Information Protection Legislation: A Comparative and Forward-Looking Perspective on the Third Anniversary of the EU GDPR's Entry into Force]. *Lexology*, May 26, 2021. <https://www.chinalawinsight.com/2021/05/articles/compliance/个人信息保护立法效果、理念及价值平衡-欧盟-gdprg/>.

"《个人信息保护法》与 GDPR 主要条文对比简析（一）" [A Brief Comparative Analysis of the Main Provisions of the PIPL and the GDPR (Part 1)]. *Lexology*, May 5, 2022. <https://www.lexology.com/library/detail.aspx?g=9da4b4b3-ca1a-40a0-a5fe-de6dc8d836fa>.

Ministry of Public Security of the People's Republic of China. *Measures for the Supervision and Inspection of Cyberspace Security by Public Security Organs* [公安机关网络空间安全监督检查办法], Order No. 176. Promulgated August 6, 2026, effective October 1, 2026. <https://www.mps.gov.cn/n6557558/c10576497/content.html>.

"NEWS: Sanders, Ocasio-Cortez Announce AI Data Center Moratorium Act." Office of Senator Bernie Sanders, March 25, 2026. <https://www.sanders.senate.gov/press-releases/news-sanders-ocasio-cortez-announce-ai-data-center-moratorium-act/>.

"NEWS: Sanders Statement on Upcoming U.S.-China AI Talks." Office of Senator Bernie Sanders, May 11, 2026. <https://www.sanders.senate.gov/press-releases/news-sanders-statement-on-upcoming-u-s-china-ai-talks/>.

Peking Nology. "Bernie Sanders Invited Two Chinese AI Researchers to Talk Safety Cooperation. Here Is What They Said." *Peking Nology* (Substack), May 9, 2026. <https://www.pekingnology.com/p/bernie-sanders-invited-two-chinese>.

"Sanders Convenes Chinese and American Computer Scientists to Warn of Threat from AI." *Broadband Breakfast*, April 30, 2026. <https://broadbandbreakfast.com/sanders-convenes-chinese-and-american-computer-scientists-to-warn-of-threat-from-ai/>.

"Sanders Urges Cooperation with China to Avert AI Apocalypse." *Responsible Statecraft*, April 30, 2026. <https://responsiblestatecraft.org/us-china-ai-race/>.

Sanders, Bernie. "The Existential Threat of AI and the Need for International Cooperation." Video, Office of Senator Bernie Sanders, April 29, 2026.

<https://www.sanders.senate.gov/videos/the-existential-threat-of-ai-and-the-need-for-international-cooperation/>.

"Tech: Republicans Call Out Sanders for AI Event." *Punchbowl News*, April 27, 2026.

<https://punchbowl.news/article/tech/republicans-sanders-ai-china/>.

"The Hill: Bernie Sanders Warns of AI's Existential Risk, Calls for US-China Cooperation." *The Hill*, May 7, 2026. <https://thehill.com/policy/technology/5866949-sanders-ai-risks-china-cooperation/>.

TISI (Taiwan Institute for Sustainable Energy, comparative law research). "中美欧个人信息保护法比较 | 附 4 万字解读报告下载" [A Comparison of Personal Information Protection Laws in China, the United States, and the EU, with a 40,000-Word Interpretive Report]. Accessed August 2026. <https://www.tisi.org/19493/>.

Xue, Lan (薛澜). "人工智能技术的社会应用——治理挑战" [The Social Application of Artificial Intelligence Technology — Governance Challenges]. Keynote address, 2026 Academic Conference on China's Digital Economy Development and Governance, Tsinghua University, Beijing, July 3, 2026. Transcript published by the Institute for Service Economy and Digital Governance, Tsinghua University, July 30, 2026.

Zhonglun Law Firm. "《个人信息保护法》与 GDPR 主要条文对比简析（一）" [A Brief Comparative Analysis of the Main Provisions of the PIPL and the GDPR (Part 1)]. *Lexology*, May 2022. <https://www.zhonglun.com/research/articles/8645.html>.

Zhonglun Law Firm. "他山之石 | 欧盟 GDPR 之镜鉴中国个人信息安全规范" [Learning from Others: The EU GDPR as a Mirror for China's Personal Information Security Specification]. May 22, 2018. <https://www.zhonglun.com/research/articles/6922.html>.

Zhonglun Law Firm. "中国和欧盟法域下标准合同合规路径比较" [A Comparison of Standard Contractual Compliance Pathways under Chinese and EU Law]. Accessed August 2026. <https://www.zhonglun.com/research/articles/9707.html>.

Zuckerberg, Mark. "The AI Future Is for Everyone." *Wall Street Journal*, July 28, 2026.

Zuckerberg, Mark. "The Future Is for Everyone." Meta, August 10, 2026.