

Platforms, Algorithms, and Discursive Power: A Comparative Audit of Top-Tier Influencer Ecosystems in China and the United States

Lihui Zhang

M-RCBG Senior Fellow at Harvard Kennedy School, Former President of Caixin Media

Abstract: Cross-national comparisons of digital influence usually reduce to follower counts. A figure such as MrBeast's 468 million YouTube subscribers set against Li Ziqi's 30 million is read as an eight-to-one ratio of cultural reach and reported as such on both sides of the Pacific. This paper argues that such numbers are compressed arguments rather than neutral facts, and often incoherent ones, because the pools they draw on, the acts they aggregate, and the legal regimes that produce them are not the same across the two systems. A raw integer stripped from its platform, its measurement protocol, and its regulatory context cannot do the comparative work it is asked to do. Drawing on a multi-platform audit of 105 top-tier accounts, 52 in the United States and 53 in China, across twelve platforms, with 628 channel-level records, the paper shows that the two ecosystems are structurally hard to place on a single scale along three axes: platform architecture, platform stack, and content regime. The empirical core is a three-axis influence index that scores each creator on political, economic, and cultural influence against a region-specific virtual ceiling, so that no real creator reaches the maximum. The paper proposes algorithmic sovereignty as a comparative frame that treats infrastructure, discursive governance, and data jurisdiction as first-order variables, and it reframes the familiar free-versus-censored binary as two sovereignty constructions with distinct trade-offs. It closes by revising the theory of the public sphere: that sphere is no longer singular but algorithmically partitioned into mutually opaque layers, with the cultural register, rather than the political one, providing what bridging still occurs.

Keywords: Influencer ecosystems, platform governance, comparative media systems, China, algorithmic sovereignty, public sphere, comparative metrics

1. Introduction

China and the United States host the two largest digital public spheres in the world, and the two most different. The gap between them has widened rather than closed over the decade since mobile internet reached saturation in both countries. In each, the leading influencers have grown into something like media institutions. They are advertisers, political actors, and agenda-setters at once, and in a few domains, consumer recommendation, political identity, and the framing of breaking news, they now reach further than the legacy press they displaced. Yet the comparative literature is thin, scattered across single-country studies, and careless about the unit of analysis.

Western platform metrics, monthly active users, subscriber counts, engagement rates, are routinely dropped into the Chinese context, or the reverse, without asking whether the underlying acts are the same. Does following a WeChat Official Account carry the same commitment as subscribing on YouTube? Does a like on Douyin signal what a like on TikTok signals, given how differently the two recommenders weight it? This is a bias problem, not a data problem. More data gathered under the same flawed assumptions would not fix the comparison. It would only make the illusion of precision sharper.

The puzzle is therefore about method before it is about evidence. When an analyst writes that MrBeast is fifteen times larger than Li Ziqi, what is being measured, and against what baseline? Subscribers on a single open platform with counts anyone can verify through an API, drawn from a pool that skews young, male, and global? Or the



combined footprint of a creator spread across six closed and semi-closed Chinese platforms whose denominators are estimated or unavailable, and whose audiences overlap to an unknown degree, so that adding them may count one viewer five times? The numbers are not wrong in a technical sense. MrBeast really does have 468 million YouTube subscribers, and the figure is auditable. They are wrong in that they answer a different question than the one they appear to answer, and the swap is usually invisible to the analyst and the reader alike.

The paper makes three contributions. It releases an open dataset of 105 top-tier creators across twelve platforms with explicit provenance tags, so later researchers can challenge or extend the claims against the same evidence rather than a proprietary black box. It introduces a measurement instrument, the three-axis influence index, that replaces the single follower number with three axis scores. And it offers a theoretical reframing, algorithmic sovereignty, through which the structural gap between the two ecosystems becomes legible without collapsing into a free-versus-censored binary that describes neither system well.

2. Research questions

Three questions organize the study. The first is structural: are there systematic differences between top-tier United States and Chinese influencers in platform structure, content category, and the public visibility of political stance? The comparison is set at the level of structure rather than individual variation, because the relevant object is two ecosystems shaped by different architectures and different rules, not the quirks of any one creator.

The second is about operationalization: how can top-tier creators in the two systems be compared with a common index that captures political, economic, and cultural influence without folding the comparison back into a single follower ranking? The instrument is part of the argument, since the form of comparison decides which structural differences become visible at all.

The third is theoretical: where do the two ecosystems meet, and where do they stay untranslatable? What does that imply for theories of the digital public sphere, for the comparative sociology of platforms, and for AI governance, especially the questions of cross-border data flow, algorithmic transparency, and the looming displacement of human creators by generative systems? The payoff is meant to travel past the China and United States pair. The frame should help with any comparison in which the platforms, the rules, and the measuring instruments are built by different sovereigns.

3. Data and method

3.1 Sample and coverage

The dataset covers 105 top-tier creators, 52 in the United States and 53 in China, mapped across twelve platforms. The United States side draws on YouTube, TikTok, Instagram, Twitter/X, Facebook, Reddit, and LinkedIn. The China side draws on Bilibili, Douyin, Weibo, Zhihu, and Xiaohongshu, with several creators, Li Ziqi most of all, appearing on both. Creators were identified from platform leaderboards, industry rankings from aggregators such as SocialBlade, Newrank, and CBNDData, and journalistic top-creator lists, with final inclusion requiring appearance on at least two independent sources to limit any one list-maker's bias. Table 1 summarizes the sample.

Attribute	United States	China
Creators	52	53
Primary platforms	YouTube, TikTok, Instagram, X, Facebook, Reddit, LinkedIn	Bilibili, Douyin, Weibo, Zhihu, Xiaohongshu
Cross-system bridge	Li Ziqi (shared)	Li Ziqi (shared)
Channel-level records	628 total across both	
Fields per record	followers, views, video count, engagement, timestamp, stable identifier	

The result is 628 channel-level records, each storing follower count, view count, video count, engagement rate, a collection timestamp, and a stable identifier that lets later passes relocate the same account even when its display name changes. Sampling is purposive, top-of-distribution by within-country stature rather than a random draw from the long tail. The boundary condition is explicit: the claims apply to the head of the distribution, not the body, and any extension to mid-tier or long-tail creators would need a different sampling strategy and would almost certainly yield different findings.

3.2 The three-axis influence index

Rather than treat influence as a single number, the index rewrites each creator as a three-number profile: political, economic, and cultural. The aim is not to infer persuasion, which the dataset cannot support, but to describe the shape of a creator's public position within the observed platform stack. The instrument is deliberately sparse and uses only quantities the dataset already holds. Table 2 sets out its components.

Component	Operationalization	Function
Reach	log of summed followers across platforms, weighted by platform	Region-specific measurable reach
Salience per axis	keyword match on category and political stance	Maps genre onto P / E / C relevance
Raw axis score	reach multiplied by axis salience	Axis-specific pre-normalization value
Reference ceiling	1.10 times the highest raw axis score in the region	Virtual ceiling, not a real creator

Each score is then the raw axis value divided by the regional reference ceiling, times one hundred. This normalization is the core design move. It abandons percentile ranking and empirical maxima in favor of a virtual ceiling that sits ten percent above the observed frontier in each region and axis. The result is deliberate headroom, so that even the strongest real creator sits below one hundred and the instrument keeps analytic distance between the sample frontier and the ceiling. Two composite summaries are retained because a single number is still sometimes useful. Peak is the maximum of the three axes and marks a creator's strongest dimension. Geom is the geometric mean of the three, which rewards balance and penalizes specialization. Read together they separate the specialist from the distributed brand.

4. Findings

The section is built around one main finding and five supporting ones. The main finding is the index itself. Once top-tier creators are written as three-number profiles rather than raw follower counts, the comparison turns structural rather than scalar. The supporting findings explain why the profiles take the shape they do across the two systems.

4.1 The index as a profile space




Figure 1. Three-axis profiles for selected frontier creators. Scores are normalized to region-specific virtual ceilings, so no creator reaches 100.

In the current sample the United States frontier is anchored by MrBeast on the economic and cultural axes and by Joe Rogan on the political axis. The Chinese frontier is split across three people: Xiao Lin Shuo on economic, Sleepy Bear News on political, and Liu Genghong on cultural. Because scores are normalized to virtual ceilings rather than empirical maxima, these are high values but not terminal ones. The score is not a disguised percentile or a restatement of who is biggest. It opens a profile space, shown in Figure 1 and detailed in Table 3. MrBeast appears as a two-axis giant on economic and cultural, Joe Rogan as a political specialist with a cultural tail, Xiao Lin Shuo as an economic specialist, and Li Ziqi as a sharply cultural case with almost no political or economic extension. The unit of comparison is no longer the person as a number but the person as a shape.

Creator	System	Political	Economic	Cultural	Profile
MrBeast	US	0.0	90.9	90.9	Two-axis giant (E, C)
Joe Rogan	US	90.9	0.0	43.3	Political specialist
Jake Shane	US	43.3	-	-	Political tail
Xiao Lin Shuo	CN	0.0	90.9	-	Economic specialist
Lei Jun	CN	-	84.9	-	Economic frontier
Sleepy Bear News	CN	90.9	-	-	Political specialist
Liu Genghong	CN	-	-	90.9	Cultural frontier
Li Ziqi	bridge	0.0	0.0	84.7	Cultural bridge

4.2 Head-of-distribution asymmetry

The United States ranking is shaped by a single super-giant. MrBeast, at roughly 468 million YouTube subscribers, sits in a class of one, with the second-ranked creator trailing by an order of magnitude, and the lead

reinforced by a feedback loop in which production budgets buy higher view counts that then fund the next production cycle. The Chinese ranking is flatter at the head. There are more creators in the upper tier but a lower individual ceiling, and the top five are separated by much narrower reach gaps than in the United States. The index does not erase this asymmetry. It renders it comparable in a form that raw subscriber arithmetic cannot. Once reach is weighted within region and normalized to a virtual ceiling, the United States head still looks steeper and the Chinese head still looks broader, and that relationship can be stated without falling back on naive follower counts.

4.3 Divergence of the platform stack

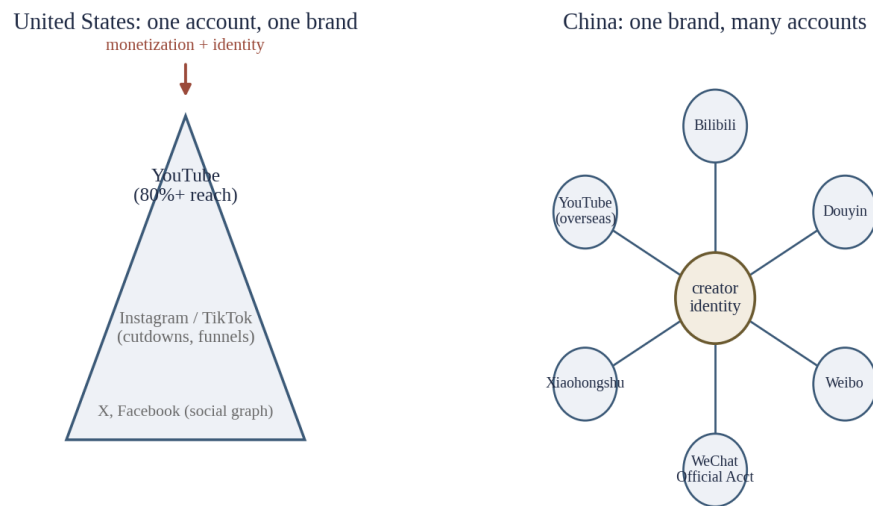


Figure 2. Platform-stack topology. The United States creator concentrates on one dominant platform; the Chinese creator distributes one brand across many accounts.

In the United States the typical top-tier creator anchors on YouTube, drawing more than eighty percent of measurable reach from that one platform. Instagram and TikTok serve as distribution channels for short cutdowns and as funnels back to the long-form channel, while Twitter/X is the social-graph layer where creators talk to one another and to journalists. The form is one account, one brand. The dominant platform is both the money surface and the identity layer, and for most purposes a creator's YouTube channel simply is the creator.

In China the typical top-tier creator is pushed onto five or six platforms at once: Bilibili plus Douyin plus Weibo plus WeChat Channels plus a WeChat Official Account, with a YouTube outpost for overseas reach and, increasingly, a Xiaohongshu presence for the younger female audience that Douyin does not fully capture. No single domestic platform dominates across age groups and content types, and platform-specific algorithmic preferences make reuse across platforms a labor-intensive job rather than a copy and paste. The form is one brand, many accounts, held together by the creator's identity and by a back-office team that adapts content, engineers audience funnels, and stitches monetization across a surface that no platform wants to unify. The two heads barely intersect. Chinese top-tier creators are essentially absent from X, Reddit, and LinkedIn even when they have the language skills and the ambition, and United States top-tier creators are essentially absent from Chinese platforms even in categories with a clear Chinese audience. Li Ziqi is the one clear exception across the whole sample, and that she is the only one is itself a structural finding: the head of the distribution is sealed by jurisdiction almost completely.

4.4 Content-category skew and political visibility

The United States top tier skews toward entertainment, lifestyle, gaming, and comedy, which together carry most of the measurable cultural mass. The Chinese top tier skews toward live commerce, knowledge, and short drama, a much higher share of transactional content in which the path from viewing to purchase is built into the format rather

than appended as sponsorship. In index terms this is why the United States distribution fills the cultural axis while the Chinese distribution populates both economic and cultural. A distinctive Chinese formation is the knowledge commentator, treated in Section 6, which has no direct equivalent in the United States top tier.

The political axis is stark on first reading: only a small subset of the sample scores above zero. Joe Rogan anchors the United States axis, with Jake Shane a distant second; Sleepy Bear News anchors the Chinese axis, followed by Luo Xiang and Mashu Bobi. This is not evidence that politics is absent elsewhere. It shows that in the current coding scheme political influence becomes legible only where category and stance metadata clear a fairly high salience threshold. The contrast is structural rather than psychological. The United States axis shows a politicized top specialist and a more visible tail; the Chinese axis concentrates in a small cluster of knowledge, law, and history. Nothing here supports a reading in which one side is less political in any civilizational sense.

4.5 The Li Ziqi case

Li Ziqi is the only creator in the sample whose comparative significance is captured better by a profile than by a platform inventory. In the current data she appears as almost pure cultural influence, with neither political nor economic extension. That profile matters because it isolates the one register that travels cleanly across the two systems. Her content is slow, heritage-coded, minimally verbal, and largely non-transactional. It succeeds not by beating the two governance filters but by staying out of the zones that trigger them. She is the exception that defines the rule. Cross-border cultural reach appears to require a narrow set of conditions: content that is minimally verbal, so translation loss is small; culturally specific but not culturally contested, so it keeps its novelty without inviting friction; and slow enough to survive compression into any short format. The narrowness of that condition set is why she remains singular.

5. The two ecosystems in profile

The United States ecosystem is a YouTube-centered pyramid, with Instagram and TikTok as secondary distribution and the older platforms reduced to social-graph and commentary roles. The pyramid is steep. The distance between the top creator and the twentieth is roughly an order of magnitude in subscribers and more than that in production budget, and the steepness is reinforced by a platform economy that rewards scale with better algorithmic treatment and better advertiser access. Monetization runs through advertising, memberships, merchandise, revenue share, and direct-to-fan commerce. Political expression is open and easy to code, with long-form podcasting having taken over the civic-commentary niche that network television vacated. Data access is a gold standard on the dominant platform, since anyone with a YouTube API key can check the claim that MrBeast has 468 million subscribers, and that very checkability is a civic good whose loss on other platforms has degraded the public's ability to reason about scale. The main risk is heavy single-platform dependence: the ecosystem's collective bet on YouTube is a bet that one company's incentives stay aligned with the creator economy, and that assumption is not guaranteed.

The Chinese ecosystem is a multi-platform mesh with no single dominant node and strong age and format loyalty that keeps any creator from consolidating an audience on one surface. The mesh is both a constraint and an insurance policy. No creator can be removed by a single platform's decision, because no creator lives mainly on one platform, but no creator can reach the vertical scale that single-platform dominance would allow. Monetization runs through live-commerce sales, platform creator funds, and private-domain conversion, the WeChat funnel that moves public-discovery audiences into direct-message relationships with far better retention economics than any public ranking permits. Political expression is heavily filtered, and the civic-commentary niche is occupied by the knowledge-commentator cluster rather than by open podcasting. Data access is bifurcated, which is the most analytically important feature of the system. Bilibili and Zhihu interfaces are usable; Douyin is intermittently hostile to outside collection; Weibo has been closing its public interfaces for years; and WeChat is effectively a black box, observable only through what creators choose to disclose. The index reflects this: China fills the economic axis more densely and still sustains a strong cultural frontier, but it does so through a measurement environment with larger blind spots.

6. The knowledge-commentator cluster

A distinctive Chinese formation recurs often enough in the sample to count as a coherent social type rather than a set of individual cases. Its members concentrate on Bilibili and WeChat, with comparatively weak Douyin performance, because the short-form recommender does not suit the cognitive density of the content and the audience self-selects toward interfaces that reward sustained attention. They occupy the civic niche that in the United States belongs to long-form podcasting, but they keep video as the primary medium. Their content runs long, typically thirty minutes or more, and it positions itself against the short entertainment that dominates the platforms it lives on. Their monetization mixes platform incentive payments, paid knowledge products, and brand deals aimed at an audience defined by professional identity, with less direct subscriber revenue than United States counterparts because the subscription tooling available to Chinese creators is thinner and the audience relationship runs through the private WeChat funnel rather than a recurring paid tier.

Representative figures span finance, in Xiao Lin Shuo, whose credential trajectory is central to the credibility he deploys; industrial policy, in Sleepy Bear News, who works in a register that is national-interest-aligned while keeping enough analytical specificity to avoid pure cheerleading; history and military affairs, in Mashu Bobi, whose comparative narratives sit at the edge of popular history and strategic commentary; and criminal law, in Luo Xiang, whose case-based analysis pairs professional standing with platform-friendly pacing. Table 4 places the Chinese and United States versions of the type side by side.

Axis	China	United States
Form	Long-form video plus illustrated text (Bilibili and Official Account)	Long-form podcast (audio and YouTube)
Niche	Fills the post-broadcast gap for serious public discussion	Same
Political visibility	Indirect, through industrial, legal, and historical framing	Direct, explicit positional commentary
Monetization	Platform incentives, knowledge payment, brand partnerships	Advertising, Patreon, direct subscription

The comparison crystallizes the paper's central claim. An identical template, the expert-commentator personal brand, is shaped by algorithm and regulation into two distinct civic forms, and the difference is neither culture nor talent but the structural conditions under which the form must be produced and paid for. The United States version can afford to be direct because the regulatory cost of a political stance is low and direct-subscription monetization rewards sharp positioning. The Chinese version is required to be indirect because the regulatory cost of a direct stance is prohibitive and the WeChat-funnel logic rewards content that can travel across professional networks without tripping platform-level risk signals. This is not a claim that Chinese creators are less aware or less capable. On several measures the analytic sophistication of the Chinese cluster runs higher than the average register of the United States podcast niche. It is a claim that identical inputs on the creator side are transformed into divergent outputs by platform architecture and regulatory environment, and that these transformations are what comparative analysis has to address [15].

7. Algorithmic sovereignty

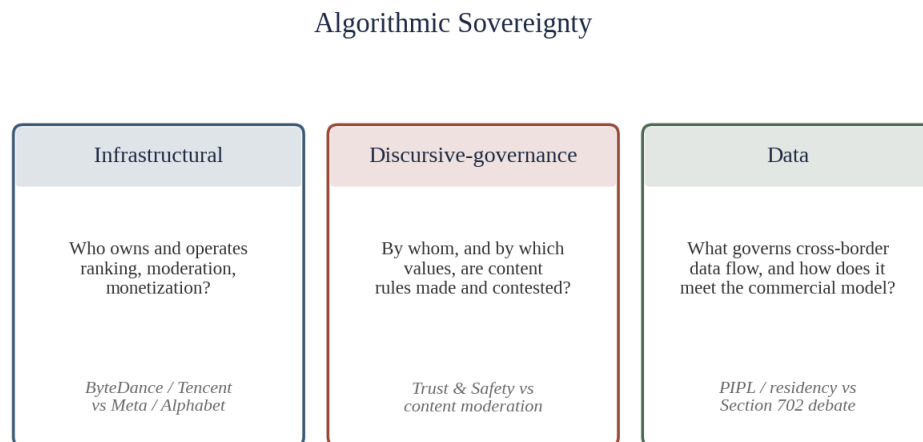


Figure 3. The three dimensions of algorithmic sovereignty.

Algorithmic sovereignty is the frame that organizes these findings into a whole. The choice of sovereignty rather than power or control is deliberate, meant to invoke the tradition in which sovereignty is a bundle of jurisdictional claims rather than a single capacity. The frame has three dimensions, shown in Figure 3, each tied to a distinct set of actors and a distinct trade-off. Infrastructural sovereignty asks who owns the platform, meaning not the equity stake alone but the operational control over ranking, moderation, and monetization that the stake nominally confers. Discursive-governance sovereignty asks by whom, and by which values, content rules are made, and by what process they can be contested. Data sovereignty asks what governs cross-border data flow, and how those constraints interact with the platform's commercial model and its regulatory regime.

The difference between the two systems is not the binary of freedom versus censorship that dominates popular framing. That framing does real damage by collapsing three dimensions into one and projecting onto Chinese platforms a uniformity they do not have. What we have instead is two different sovereignty constructions, each with a distinct trade-off and characteristic vulnerabilities. The United States construction maximizes infrastructural openness and discursive permissiveness while staying indifferent to data sovereignty until very recently. The debate over TikTok divestiture is the moment that indifference became unsustainable, and the resulting improvisation showed how unprepared the United States regulatory apparatus was to defend a dimension it had never theorized as such. The Chinese construction prioritizes discursive and data sovereignty at the price of infrastructural plurality, accepting a compressed expressive space and a crowd of platforms as the cost of keeping domestic control over both rule-making and data residency [12], [13], [14].

8. Discussion: a partitioned public sphere

Classical theory assumed a single rational-critical arena in which participants set aside private interest and argued under shared norms, with the better argument expected to prevail over time. The 105-creator, twelve-platform dataset shows something else. The public sphere has been partitioned algorithmically into mutually opaque layers, and the partition has gone far enough that the singular arena is no longer even a useful fiction against which to measure platform reality. Both countries show this partition. The logic of the partition differs, and that difference is the pivot of the comparison.

In the United States the partition is produced by engagement-optimizing ranking on nominally open platforms, loosely bridged by spillover. A TikTok clip migrates to YouTube, then to podcast discussion, then back to TikTok,

and the channels between platforms carry both content and audience in a porous if uneven way. In China the partition is produced by ranking and regulation together, with hard isolation between platforms that blocks the cross-platform meme migration typical of the United States system, and a moderation regime that actively polices the borders of each platform's permissible-speech envelope. The theoretical implication is direct. The single rational-critical public has to be replaced by affective publics in the plural, spheres partitioned by protocol and glued together, where they are glued at all, by the cultural rather than the political register. The Li Ziqi case is the empirical illustration: the only stable bridge across the partition is a register that is explicitly non-political, and that is a systemic property of the partitioned sphere rather than a fact about one creator. This reading extends a thread that runs through the public-sphere tradition and its critics, from the original account of a bourgeois arena to later work on affective publics, on platforms as governors rather than conduits, and on networked expression under different speech regimes [1], [2], [3], [6], [7], [8].

9. Implications and limitations

For researchers, the practical implication is to stop comparing raw Chinese and United States follower counts as if they were already commensurable, and to treat a multi-axis instrument as the baseline expectation for comparative claims, with index refreshes a normal part of the work rather than an embarrassment. For policy, debates over TikTok divestiture, data-protection law, and content rules rest on implicit quantitative assumptions that usually take the form of appeals to scale [12]. Making those assumptions explicit tends to narrow the claims that can be supported and sometimes reverses them. One testable hypothesis follows: generative systems will hit top-tier creators before they hit ordinary users, and the first wave of displacement will concentrate in the Chinese live-commerce segment, where the host persona is most replaceable and the dependence on algorithmic recommendation is highest [13].

The study has clear limits. The sample of 105 is purposive, not random. It characterizes the head, not the long tail, and any extension to mid-tier or long-tail creators runs against the paper's own disclaimer [14]. The sampling frame favors creators whose reach is public enough to be ranked, which undercounts private-domain-concentrated Chinese creators. The data is a time-slice from the first quarter of 2026, not a longitudinal panel, and the ecosystem shifts fast enough that concrete rankings will move even if the structural findings hold [15]. The index itself has boundaries: its salience values come from category and stance text, its ceiling is an operational anchor rather than a natural constant, and its scores are tied to the current snapshot [16]. Content coding remains partly manual. Most importantly, the dataset contains no audience-side data. What is measured is creator output, not reception or persuasion. This is a structured index of visible influence-shape, not influence in the strong political-science sense [17].

10. Conclusion

The two ecosystems place opposite structural pressures on the creator, and that opposition, rather than any simple ranking of freedom or capacity, is what the comparison should foreground. The United States combines winner-take-all dynamics with single-platform depth, which produces very high individual ceilings paired with platform-level fragility. China combines multi-platform dispersion with mandatory algorithmic adaptation, which caps individual ceilings but produces greater ecosystem-level resilience. Everything downstream, the advertising-versus-commerce split in monetization, linear versus adaptive content reuse, explicit versus oblique political risk, and differing exposure to generative displacement, follows from that single structural difference and is a corollary of the comparison rather than an independent finding. To compare the two ecosystems only at the level of creator behavior, as most of the existing literature does, is to compare the surface expression of two different constraints without naming the constraints. The three-axis index and the algorithmic-sovereignty frame are offered as one way to name them, and to keep a stable comparative language as the data is refreshed.

REFERENCES

1. Habermas, J. (1989). *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society* (T. Burger, Trans.). Cambridge, MA: MIT Press. (Original work published 1962).
2. Papacharissi, Z. (2015). *Affective Publics: Sentiment, Technology, and Politics*. New York: Oxford University Press.

3. Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. New Haven: Yale University Press.
4. van Dijck, J., Poell, T., & de Waal, M. (2018). *The Platform Society: Public Values in a Connective World*. New York: Oxford University Press.
5. Porter, T. M. (1995). *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton: Princeton University Press.
6. MacKinnon, R. (2012). *Consent of the Networked: The Worldwide Struggle for Internet Freedom*. New York: Basic Books.
7. Benkler, Y. (2006). *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. New Haven: Yale University Press.
8. Yang, G. (2009). *The Power of the Internet in China: Citizen Activism Online*. New York: Columbia University Press.
9. Zhao, Y. (2008). *Communication in China: Political Economy, Power, and Conflict*. Lanham, MD: Rowman & Littlefield.
10. Poell, T., Nieborg, D., & van Dijck, J. (2019). Platformisation. *Internet Policy Review*, 8(4), 1-13.
11. Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.
12. Ankam, J. (2024). "Contracts Across Modalities: Detecting Embedding Staleness and Governance Drift in Multimodal Lakehouse Architectures." *International Journal of Computational and Experimental Science and Engineering*, 10(1). <https://doi.org/10.22399/ijcesen.5469>
13. Ankam, J. (2022). "Benchmarking Autonomy: A Taxonomy of Human-in-the-Loop Checkpoints for AI-Generated Transformation Code." *International Journal of Computational and Experimental Science and Engineering*, 8(3). <https://doi.org/10.22399/ijcesen.5462>
14. Vytla, S. K. (2025). "DEVOPS-GRID: DevOps-driven reliability and digital-twin operations for intelligent power systems." *GongchengKexueYu Jishu/Advanced Engineering Science*, 57(3).
15. Vytla, S. K. (2026). "GOVERN-CTRL: A Governance Control Plane Architecture for Production-Scale Data and Machine Learning Pipelines." *International Journal of Computational and Experimental Science and Engineering*, 12(3). <https://doi.org/10.22399/ijcesen.5394>
16. Vytla, S. K. (2024). "POLICY-ML: Policy-as-Code Enforcement for Compliance, Auditability, and Trust Across Data and Machine Learning Pipelines." *International Journal of Computational and Experimental Science and Engineering*, 10(4). <https://doi.org/10.22399/ijcesen.5393>
17. Gaggar, H. (2026). "MemGov: Policy-Governed Memory Lifecycle Management for Enterprise LLM Agents." *Journal of Intelligent Decision Making and Information Science*, 3(11s), pp. 2199–2214.