Regulating the Traffic Economy: Salary Caps, Superstar Rents, and Platform Production

Yu (Jasmine) Hao
Faculty of Business and Economics
The University of Hong Kong
yhao@hku.hk

Yele Ma
Faculty of Business and Economics
The University of Hong Kong
yelema@connect.hku.hk

Yanhui Wu
Economics; Management and Strategy
HKU Business School
The University of Hong Kong
yanhuiwu@hku.hk

Ke Wang
iQIYI, Inc.

July 2, 2026

Abstract

Platform markets can fail upstream. When platforms monetize attention through advertising, subscriptions, and subscriber acquisition, a downstream monetization metric is transmitted backward into upstream production: platforms bid for scarce inputs that predict attention, and competition capitalizes those attention rents into superstar wages. The resulting input tournament can narrow the production frontier toward high-budget, traffic-star projects and away from director-led or experimental variety. We study this mechanism using China's 2018 streaming salary cap, a public commitment device that weakened bidding for traffic stars. On 3,822 dramas, we estimate a structural model of demand, actor careers, matching, and platform project control. The cap restores positive platform bargaining power ($\widehat{\omega} = 0.599$), reveals a high pre-cap Bertrand outside-option intensity ($\widehat{\varphi} = 0.998$), and admits entrants. Counterfactual accounting implies RMB 181 billion in actor wage compression, with production responses operating through project control and director entry.

Keywords: media platforms, upstream input markets, salary caps, vertical integration, superstar tournament, variety, China streaming.

JEL classification: L82, L13, L51, J31, D43.

1 Introduction

This paper asks how platform monetization shapes upstream production. Platforms are usually studied through downstream prices, commissions, ranking, recommendation, self-preferencing, and market power. We study a different margin: when platforms monetize attention, the metric they use to earn revenue can be transmitted backward into the contracts for scarce upstream inputs. In talent-intensive media, platforms compete for traffic stars because stars predict attention, advertising demand, and subscriber acquisition. The equilibrium distortion is therefore not only that stars are expensive. It is that a whole production technology can become organized around the input that best converts into the platform's monetization metric.

China's 2018 streaming salary cap provides a setting in which this upstream platform distortion can be observed. If the cap were only a labor-cost rule, we should see lower actor pay with little change in platform control, director allocation, or the quality distribution of dramas. The data instead show a coordinated transition: platform residuals recover, production shifts toward platform-controlled projects, previously excluded directors are absorbed, actor career margins move, and the market selects into fewer but higher-scoring projects. The structural model is built to explain this joint movement rather than any one fact in isolation.

A long literature on media markets establishes that ad-funded entry under-supplies variety relative to first-best when consumer tastes are heterogeneous (Spence, 1976; Mankiw and Whinston, 1986; Anderson and Coate, 2005; Berry and Waldfogel, 1999; Waldfogel, 2003; Aguiar and Waldfogel, 2018). The mechanism is a *preference externality*: producers maximize aggregate viewership, not utility-weighted attention, so minority tastes are systematically underweighted. In the language of recent work on technology and platform incentives, advertising turns a multi-dimensional social objective into a single attention-based reward, directing innovation and production toward what is most monetizable rather than what is socially most valuable (Acemoglu and Johnson, 2023). This article argues that in talent-intensive platform markets, the preference externality is amplified by an upstream Bertrand auction for scarce stars whose welfare consequences can be disciplined empirically. We show this using the market transition around China's 2018 actor salary cap in the world's largest streaming market by audience size.

The mechanism in three steps. First, advertising and subscriber acquisition compress heterogeneous tastes into a monetizable attention metric. Second, Bertrand competition for scarce traffic stars capitalizes that attention reward into the actor's outside option, pushing wages toward the star's full value-added: $W \simeq V_a$. Third, the inflated input price compresses the platform residual that would otherwise finance platform-controlled production capability and entry by non-incumbent directors. The cap works as a public commitment device because it weakens the auction over the same attention asset that private platforms could not credibly stop bidding for.

The 2018 cap as commitment device. In September 2018, Chinese regulators imposed binding restrictions on actor compensation, subsequently clarified to require that total actor pay not exceed 40% of total production cost (C1) and that lead-actor pay not exceed 70% of total cast compensation (C2). Together these bind top-star compensation to 28 percent of production cost, against the 50–70 percent share top stars routinely absorbed pre-cap. The three major platforms had attempted a private non-bidding-war agreement in mid-2018 that failed enforcement within months—consistent with a coordination failure (Cooper et al., 1990; Greenwald and Stiglitz, 1986) resolved only by public regulation. The cap is best read not as wage control in the narrow sense, but as a *public commitment device* that achieves what private antitrust cannot: it breaks the Bertrand auction by removing the star's outside option in streaming, restores positive platform surplus, and unlocks the variety-feasible production technologies that were budget-foreclosed before.

What we estimate. We first write a compact identification model that separates the salary-cap commitment channel from contemporaneous industry shocks. The identifying object is not a single post-2018 change, but the joint movement of platform residual surplus, project-control rights, director entry, and market selection after a nationwide cap that followed a failed private non-bidding agreement. We then build a structural model that takes the consumer-demand side from Berry et al. (1995) and Nevo (2001) BLP random-coefficients, the talent-market wage formation from Nash bargaining (Ho and Lee, 2017), and a project-level endogenous control-rights regime that closes each platform's producer problem. The model is organized in three layers. The value-creation layer estimates drama demand and star marginal revenue product. The surplus-and-control layer maps that value into wages, platform residual surplus, and platform-specific make-or-buy decisions. The career-allocation layer separates the extensive margin of lead-offer arrival from the intensive margin of actor–director sorting and actor acceptance. The core producer-transition primitives are the BLP preference coefficients, the Nash bargaining weight ω, the Bertrand outside-option intensity φ, platform-specific control-rights cost wedges $\Delta\kappa_{0p}$, the capability cost-slope a_κ, and the post-cap capability-decay rate λ. We estimate them sequentially on a panel of 3,822 streaming dramas (2017–2025), using disclosed iQIYI project-cost and revenue moments to calibrate ω, the pre-cap residual moment plus three-platform project-control shares for iQIYI, Tencent, and Youku to estimate φ, a_κ, λ, and $\Delta\kappa_{0p}$, and the actor career-path Bellman from a Hotz-Miller two-step CCP procedure (Hotz and Miller, 1993; Arcidiacono and Miller, 2011).

Headline estimates and their welfare implications are reported in the abstract and developed in detail in Sections 5 and 7.

Contributions and connection to the literature. The main contribution is to show that platform distortions need not arise only from downstream market power or algorithmic steering. When platforms monetize attention, competition for scarce upstream inputs can convert attention rents into input rents and foreclose whole classes of products. Input-price regulation can therefore operate as platform regulation by relaxing the production constraint that sustains monoculture. We connect the preference-externality literature in media markets (Spence, 1976; Mankiw and Whinston, 1986; Anderson and Coate, 2005; Berry and Waldfogel, 1999) to superstar-rent models (Rosen, 1981; Adler, 1985; Frank and Cook, 1995) and platform competition (Rochet and Tirole, 2003; Armstrong, 2006). Empirically, we assemble a 3,822-drama China streaming panel around a binding salary-cap shock. Structurally, we decompose the policy into talent bargaining, platform-specific project-control decisions, persistent control capability, and director entry, estimating ω, φ, $\Delta\kappa_{0p}$, a_κ, and λ as the primitives that govern the transition.

Roadmap. Section 2 gives the institutional background and data construction. Section 3 presents descriptive evidence on the equilibrium transition. Section 4 combines the structural model with the identifying moments for each block; Section 5 reports estimates and fit; Section 6 interprets the estimates; Section 7 reports counterfactuals; Section 8 concludes. Detailed derivations, parameter tables, and robustness checks are in the appendices.

2 Background

This section asks what institutional and data variation can make the policy interpretable as a market-design shock rather than as an accounting rule. If the cap applied narrowly or the data observed only outcomes, the paper could not separate wage compression from production reallocation. The institutional record says the cap was broad, enforceable through licensing, and preceded by a failed private platform agreement; the data combine drama outcomes, person-level careers, platform accounts, internal project-control labels, advertising records, and endorsement records. That combination lets us follow the policy from platform residuals to production control, talent allocation, and market output.

2.1 Industry background

The industry question is why actor compensation could shape production organization in the first place. If Chinese streaming platforms were passive distributors buying finished content in a thick spot market, a salary cap would mainly affect upstream studios. The institutional setting is different. Long-form video streaming is a concentrated oligopoly: iQIYI (Baidu), Tencent Video, and Youku (Alibaba) account for the bulk of paid subscribers and content spend, with Mango TV (Hunan Broadcasting) as a smaller fourth platform included in the demand controls and market-context evidence. These platforms do not only exhibit completed dramas; they commission, finance, acquire, promote, and increasingly control production. The relevant market is therefore a platform-mediated production market, not a pure downstream exhibition market.

The revenue model makes attention the central reward. Platforms monetize dramas through a mix of subscription conversion, advertising inventory, sponsorship exposure, and cross-promotion inside platform ecosystems. In iQIYI's audited Form 20-F filings, online advertising revenue alone accounts for 45 percent of total revenue on average during 2016–2018, before the subscription pivot becomes dominant (iQIYI, Inc., 2018–2024). A drama with a traffic star can raise expected attention before release, improve advertiser demand, and help platforms compete for subscribers. If attention were diffuse across titles, this would not create severe upstream bidding pressure. But the advertising and viewership data show a tournament structure: a small number of actors and titles capture a disproportionate share of attention, while the long tail receives much thinner sponsorship. This is the industry background for the bargaining mechanism in the model: platforms bid for scarce traffic stars because stars are a privately valuable way to buy expected attention. Appendix Table 26 collects the external evidence behind this interpretation.

Production organization also matters. Before the cap, platforms relied heavily on licensed or externally controlled projects, while external studios and adjacent media—broadcast television and theatrical film—supplied much of the experienced production talent. After the cap, audited iQIYI accounts show a shift away from licensed content and toward self-produced or commissioned content. We interpret this as a shift in control rights: platform controlled includes pure self-production and commissioned/custom production where the platform controls budget, contracting, and project direction; externally controlled includes licensed, bought, and jointly procured content where an outside studio primarily organizes production before platform acquisition.

The salary cap arrived during a broader attention-market squeeze. Short-form video, especially ByteDance's Douyin, expanded domestic daily active users from approximately 30 million in 2017 to more than 750 million by 2022, placing parallel pressure on the long-form attention budget. If this short-form shock alone explained the observed transition, we would expect a general contraction in long-form output without a sharp reorganization of project control or talent allocation. The data instead show both forces: total tracked drama supply becomes more selected, but platform residuals recover and platform-controlled production rises. Appendix Figure 14 shows the attention-market time series, and Appendix Table 28 reports advertising, pipeline, and entry-margin diagnostics.

The 2018 salary-cap policy. On 11 September 2018, the State Administration of Radio and Television and the All-China Federation of Radio and Television issued "Opinions on Further Strengthening the Management of Television Drama and Online Audio-Visual Programs." A November 2018 clarification translated the announcement into two hard ceilings binding on all streaming-drama productions:

(C1) $\sum_{a \in A_j} W_a \leq 0.40 \cdot C_j$ (total-cast cap: at most 40% of production cost).
(C2) $\max_{a \in A_j} W_a \leq 0.70 \cdot \sum_{a' \in A_j} W_{a'}$ (lead-actor cap: at most 70% of total cast compensation).

Together these bind top-star compensation to $0.40 \cdot 0.70 = 28$ percent of production cost, against the 50–70 percent share top stars routinely absorbed pre-cap. The policy applied to commissioned productions and acquired licenses alike, enforced through SART's broadcast-license approval channel. We use the September 2018 announcement as the policy cutoff. For aggregate platform-year moments, the post-cap indicator equals one from 2019 onward to accommodate the typical 12–18-month production cycle. For project-level interpretation, however, release year is an imperfect proxy for exposure because some 2019 releases had completed filming before the policy. We therefore audit public production-timing records and separate clean post-cap production starts from grandfathered projects whenever the timing evidence is used directly.

2.2 Data sources

The first data source is a drama-level performance panel that links Yunhe viewership records to Douban title records. The unit of observation is a Chinese-language drama series. Coverage is 2017–2025, which spans the pre-cap years, the 2018 policy transition, and the post-cap streaming market. The variables include release year, platform, genre, episode count, director, top-five billed cast, Douban score, rating availability, and standardized 30-day views.

The second data source is a Douban filmography panel for directors and actors. The unit of observation is a person-title credit, which we aggregate to person-year histories. Coverage is 1997–2026 across film, television, and online productions. The variables include credit year, medium, role, director identity, actor identity, billing order, debut year, annual work counts, and the rating class of past works.

The third data source is platform-side accounting and project data from iQIYI. It combines public Form 20-F filings with an internal project-level panel. The public filings cover 2016–2022 at the firm-year level; the internal panel covers 2015–2030 at the drama-project level. The variables include annual revenue, cost of revenue, subscriber counts, operating income, revenue decomposition, per-drama cost components, per-drama revenue components, acquisition route, involved parties, and production-control status. For the three-platform make-or-buy block, these iQIYI control labels are combined with the cross-platform drama panel to construct platform-year project-control shares for iQIYI, Tencent, and Youku; cells without direct sourcing labels are completed by transparent classification rules and reported as partly imputed.

The fourth data source is Endata industry data. It consists of three linked panels: brand sponsorship, television cast, and actor endorsement records. The brand-sponsorship panel covers 2017–2023 at the drama-quarter level; the television cast corpus covers 1970–2024 at the title-actor level; and the endorsement panel covers 2016–2024 at the actor-year level. The variables include brand counts, brand-exposure time, starring actors, actor rank, cast-list histories, endorsement contracts, and endorsement counts.

The final data sources provide external validation and market context. iQIYI pipeline announcements are project-announcement records covering 2015–2025; Netflix Form 10-K filings are firm-year accounting records covering 2014–2022; and the Douyin series records domestic daily active users for 2017–2022. The variables include announced project counts, lead-actor mentions, platform-controlled versus acquired status, Netflix revenue and cost of revenue, and Douyin domestic DAU.

2.3 Sample construction and treatment variables

Drama-level sample. The estimation sample is the 3,822 Chinese-language drama series in the Yunhe-Douban panel with release year in $\{2017, \ldots, 2025\}$. Yunhe Data provides a standardized 30-day viewership measure comparable across platforms; we use $\log V_{30}$ as the audience-side demand outcome. Douban provides user scores on a 1–10 scale ($N = 2{,}394$ rated dramas) and a binned rating-class proxy for unrated dramas. Cast information is restricted to the top-five billed actors per drama; billing order is recorded.

Director cohort. For the director-entry DiD, we construct a 1,045-director cohort from the Douban filmography panel, selecting individuals with at least two directing credits recorded for year before 2017.

Actor value added. Actor-level value added is the billing-order-weighted average of year-demeaned Douban scores (rating-VA) and log 30-day viewership (views-VA) from the actor's prior works. The pre-2018 maximum over the top-five billed actors enters BLP demand and the actor career model.

Treatment timing audit. The main aggregate moments use release-year exposure because platform accounts, control shares, and director histories are observed at annual frequency. This is appropriate for the broad regime shift, but it is too coarse for project-level identification. To measure the likely direction of the error, we collected public production-timing records for 19 priority iQIYI projects. Among the 16 titles released after the September 2018 cutoff, seven had already finished filming before the policy, five started filming after the cap, one straddled the cutoff, and three lacked usable filming dates. Table 1 reports the audit. The implication is conservative: release-year post indicators are used for aggregate descriptive moments, while project-level timing claims should either use the clean post-cap filming-start group, control for grandfathered titles, or exclude titles with missing production dates. Because actor contract dates are not observed in the current data, the audit does not identify contract-level exposure; it only reduces treatment error relative to release-year timing.

Table 1: External Production-Timing Audit

Audit category	Count	Share	Interpretation
Priority projects with public timing record	19	100.0%	External audit sample for treatment timing.
Released after September 2018 cutoff	16	84.2%	Release-year post exposure used in aggregate moments.
Grandfathered: filmed before cutoff	7	43.8%	Post-release titles that should not be treated as clean post-cap production.
Post-cap filming start	5	31.2%	Cleaner treated group for project-level timing evidence.
Filming window straddles cutoff	1	6.2%	Partial exposure; best used as a separate or robustness category.
Post-release title with missing filming dates	3	18.8%	Excluded from clean timing tests unless additional sources are found.
Month/year-level filming precision	5	26.3%	Useful for broad classification, not for day-level regression discontinuity.
Observed actor contract dates	0	0.0%	Contract timing remains unobserved and is not inferred from filming dates.

The audit uses public production-timing records for 19 priority iQIYI projects collected from external source pages. The cutoff is the September 2018 salary-cap announcement. Shares for timing classes are relative to the 16 titles released after the cutoff; the month/year precision share is relative to all audited projects. The table motivates using release-year post indicators for aggregate platform-year moments, while treating project-level release-year post as measured with error unless filming or contract timing is observed.

Summary statistics. Appendix Table 24 reports summary statistics for drama outcomes, actor work performance, and collaboration quality. Missing Douban ratings are not random—a score is typically shown only after at least 500 user ratings, so missing values concentrate in the lower-exposure tail. We therefore interpret Douban-based results with caution and complement them with view-based outcomes plus controls for exposure.

3 Descriptive Evidence

This section asks what changed in the data after the cap. If the policy were only a wage ceiling, the direct prediction would be lower actor compensation and little systematic movement in the organization of production. The data do not look like that. The cap arrives after iQIYI's content-cost-to-revenue ratio rises from 101.8% in 2016 to 108.6% in 2018 and its implied content-margin proxy falls from −1.8% to −8.6%; after the cap, residual surplus recovers, project control moves toward platforms, incumbent streaming directors retreat while outside directors enter, actor career margins shift, and market output becomes more selected.

Each fact below has a specific role in the structural exercise. Platform residuals discipline the bargaining regime; the control-share paths discipline the project-control transition; director reallocation pins down the producer-entry margin; actor allocation and endorsements validate the actor outside-option channel; and quantity/quality changes provide held-out validation on market selection. Table 2 collects the reduced-form evidence used by the identification model before the structural estimation assigns counterfactual magnitudes.

3.1 Platforms: residual surplus recovers

The first effect is on platform residual surplus. Before the policy, iQIYI's content costs exceeded content revenue, and the firm-level residual proxy $\widehat{\gamma}_{p,t}$ deteriorated through 2018. After the cap, the same accounting series reverses sharply: the residual proxy rises to +23.0% by 2022. This fact motivates the bargaining block: the model must explain why the cap coincides with a recovery of platform residual rather than merely lowering an accounting cost. The evidence should be read with two qualifications. First, the underlying revenue and cost-of-revenue items come from audited Form 20-F filings, so the accounting series is reliable. Second, $\widehat{\gamma}_{p,t}$ is a residual proxy, not a direct wage or bargaining-share observation: cost of revenue includes actor compensation, production costs, licensed-content amortization, and other operating costs. Because the annual series is short, we use it as a disciplined moment and mechanism check, not as a standalone causal reduced-form estimate. Figure 1 and Table 3 report the underlying series.

Table 2: Identification Evidence: Commitment Channel and Main Threat

Block	Estimate	Magnitude	Interpretation
Platform residual	Post-cap mean shift	-3.5% to 8.6% (+12.1 pp)	iQIYI residual surplus recovers after the cap.
Audited project control	iQIYI self-produced share	12.3% to 36.2% (+23.9 pp)	Audited content-cost flows shift toward platform control.
Three-platform control	Post-cap control-share shift	Tencent +23.5 pp; Youku +14.8 pp; iQIYI +23.4 pp	Project-control rise appears for iQIYI, Tencent, and Youku.
Director cohort reallocation	Locked-out x post	locked-out 0% to 9.0-14.5%; incumbents 100% to 32.6-54.3%	Experienced directors absent from the 2018 streaming core are partially absorbed post-cap.
Selection validation	Platform-control score premium	0.117 Douban points; +0.884 log views	Platform-controlled projects have higher observed quality conditional on year, platform, and genre.
Attention-shock threat	Long-form attention index	1.00 to 0.44 as Douyin DAU rises 30m to 750m	Short-video growth explains attention pressure, a competing force that predicts contraction rather than residual/control recovery.

Notes. This table collects the reduced-form evidence behind the identification model. It reports magnitudes rather than treating each row as a standalone causal regression. The attention-shock row is a threat diagnostic rather than a cap effect. Regression details and fragile specifications are reported in Appendix Table 20.

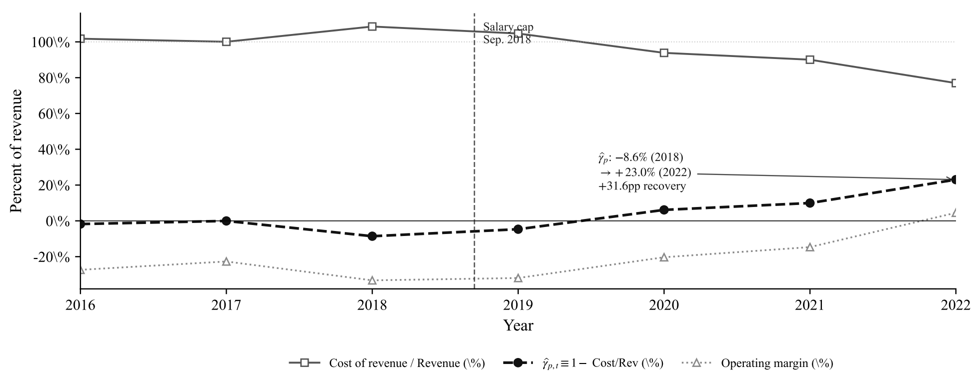

Figure 1: iQIYI content-cost-to-revenue ratio, $\widehat{\gamma}_{p,t}$, and operating margin, 2016–2022. The vertical dashed line marks the September 2018 cap announcement.

Table 3: Platform-Margin Decomposition

Year	Revenue (RMB M)	Cost of rev. (RMB M)	Cost / Rev. (%)	Op. income (RMB M)	Op. margin (%)	$\widehat{\gamma}_p$ (%)
2016	11,238	11,437	101.8	-3,074	-27.4	-1.8
2017	17,378	17,387	100.0	-3,944	-22.7	-0.0
2018	24,989	27,133	108.6	-8,306	-33.2	-8.6
2019	28,994	30,348	104.7	-9,258	-31.9	-4.7
2020	29,707	27,884	93.9	-6,040	-20.3	+6.1
2021	30,554	27,514	90.0	-4,479	-14.7	+10.0
2022	28,998	22,319	77.0	1,312	+4.5	+23.0

Revenue and cost of revenues from iQIYI Inc. Form 20-F filings, SEC EDGAR. Operating income is total revenues minus operating costs (cost of revenues + SG&A + R&D). $\widehat{\gamma}_p = 1-$ Cost/Revenue is the empirical analog of the platform's bargaining share in the model of Section 4 (Platform subsection): the residual share of revenue accruing to the platform after content costs (including talent compensation).

Reading: $\widehat{\gamma}_p < 0$ in 2016–2019 indicates that content costs alone consumed more than revenue brought in, leaving no platform residual—the literal manifestation of the auction- driven $W_a \rightarrow M\Delta s_j \cdot ARPU_{adv}$ result of Section 4. The 2018 ratio of 108.6% is the peak of the bidding-war intensity; the cap was announced in September of that year. After 2018, $\widehat{\gamma}_p$ rises monotonically to +23.0% by 2022, consistent with the platform recapturing surplus that talent had been extracting in equilibrium.

3.2 Production mode: content sourcing shifts toward platform control

The second effect is on project control. iQIYI's audited content-cost flows shift away from licensed content and toward self-produced content after the cap. The self-produced share of content cost is only 9–16% in 2016–2018, rises steadily after the cap, and reaches roughly 44–48% by 2023–2024. We interpret this as a shift in control rights: the platform moves from acquiring externally controlled content toward projects it controls directly or through commissioned production. This evidence is stronger than a pure accounting margin because it describes the organization of production itself, but it should still be read as a dynamic transition rather than as a one-year causal jump. In short annual panels, post indicators and linear trends are hard to separate; the point is that the post-cap path is a sustained ramp-up in platform control, and the same pattern appears in the three-platform control-share panel. Figure 2 provides the audited iQIYI anchor for this transition; the structural block in Section 5.3 targets platform-year control shares for iQIYI, Tencent, and Youku, together with the pre-cap residual moment, to estimate φ, $\Delta\kappa_{0p}$, a_κ, and λ. Appendix Table 24 shows that, after completing missing project labels with transparent classification rules, platform-controlled projects have at least comparable observed quality and substantially higher attention.

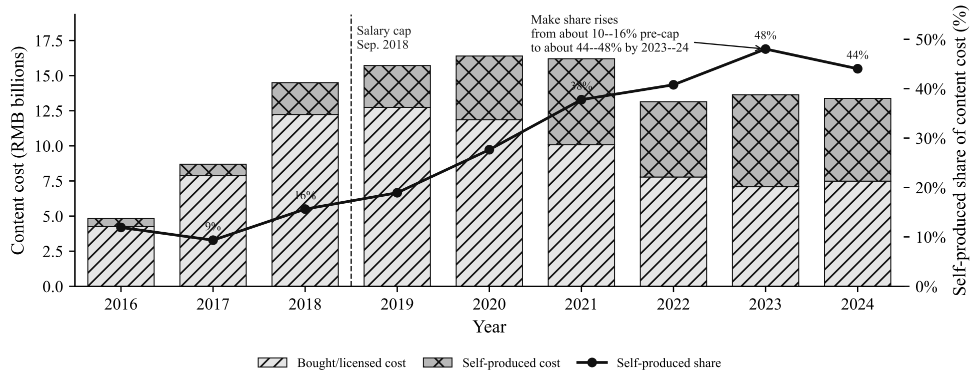

Figure 2: Dynamic project-control transition, 2016–2024. Bars report audited iQIYI Form 20-F content-cost flows for licensed/bought and self-produced content (iQIYI, Inc., 2018–2024). The line reports the self-produced share of content cost. This audited series motivates the control-rights transition; the structural make-or-buy block below targets platform-year project-control shares for iQIYI, Tencent, and Youku. The vertical dashed line marks the September 2018 cap announcement.

Table 4: Quality by Project-Control Category

	Externally controlled	Platform controlled
Matched dramas	234	479
Mean Douban score	6.09	6.16
Median Douban score	6.10	6.10
Share score ≥ 7	19.7%	21.3%
Mean log 30-day views	15.51	15.61
Controlled score gap	+0.154, $p = 0.133$	
Controlled view gap	+0.080, $p = 0.377$	

Exact-title matched internal project records to the Yunhe–Douban panel. Platform controlled includes pure self-production and commissioned/custom production; externally controlled includes licensed, bought, jointly procured, swapped-in, and distributed-in titles. Controlled gaps include platform, release-year, and genre fixed effects. The full table is Appendix Table 24.

3.3 Producers and directors: incumbents retreat, outsiders are absorbed

The third effect is director reallocation. The policy transition is not simply a story of new career entry. The visually dominant movement is a retreat of 2018 streaming incumbents: their annual streaming activity falls sharply after the cap. At the same time, the post-cap streaming market absorbs a subset of a previously locked-out outside cohort. This absorption is a level-and-count margin, not a visually strong upward trend. We therefore use "entry" in a platform-market sense: entry into the streaming-drama core, not first-time entry into directing. These outsiders are not marginal newcomers: the median locked-out director first appears in the long-run director filmography in 2004, has seven pre-2017 directing credits, and is more likely than the 2018 streaming-incumbent cohort to fall in the high-quality tier. Nor is locked-out status random. The cohort differs in observable production type: it contains fewer period-drama directors and more directors whose pre-policy work falls outside the streaming core's modal genre categories. Some members of the outside cohort are also naturally declining in all-media directing activity, so the claim is not that the entire cohort becomes active because of the cap. The interpretation is narrower: among pre-credentialed directors outside the 2018 streaming core, the post-cap market absorbs an active and available subset that had not been used in the pre-cap streaming equilibrium. Figure 3 should be read as a reallocation diagnostic. The left panel shows the visually sharp margin: incumbents were highly active before the cap and then retreat. It is not a parallel-trend test; it shows that the outside cohort was not mechanically rising before the cap. The right panel translates the same episode into director-year accounting for 2019–2022: outside-cohort absorption adds 482 director-year appearances relative to zero baseline, while incumbent activity loses 204 appearances relative to repeating the 2018 incumbent baseline. Table 5 then reports the annual rates and head counts behind this accounting: incumbent annual activity falls from a mechanical 100% baseline to 44.6% per post-cap year, while 307 of 953 locked-out directors appear in streaming at least once during 2019–2022. Because the main streaming panel has a single clean pre-period year, the long-run activity panel is a cohort-identity and incumbent-exit check rather than a streaming-entry pre-trend.

3.4 Actors: star scarcity and reallocation

The fourth effect is on actors, with two distinct margins. The first question is whether a star tournament existed before the cap. "Star tournament" is our economic label, not an industry term. The official and press record does not observe an auction directly, but it does document the two background conditions required by the mechanism: unusually high star compensation was viewed as distorting production-cost allocation, and platform revenue rewarded concentrated attention. In 2018, the Fan Bingbing tax case was linked

Table 5: Director Reallocation and Entry Margin

Panel A: Locked-out cohort DiD

	(1) Year FE	(2) + Director FE	(3) Established (5+)	(4) Weighted
Locked-out × Post	0.681***	0.681***	0.712***	0.719***
	(0.030)	(0.033)	(0.037)	(0.041)
Year FE	Yes	Yes	Yes	Yes
Director FE	No	Yes	Yes	Yes
Observations	5,225	5,225	3,560	5,225
R^2	0.195	0.452	0.441	0.434

Panel B: Cohort profile

Cohort	N	First credit	Pre-2017 credits	High quality (%)	Other genre (%)
Locked-out cohort	953	2004	7	38.1	22.5
2018 streaming incumbents	92	2003	11	23.9	1.1

Panel C: Entry and exit margins behind the DiD

	Directors	Active in 2018	Avg. active, 2019–2022	Ever active, 2019–2022
2018 streaming incumbents	92	100.0%	44.6%	82
Locked-out cohort	953	0.0%	12.6%	307

Panel D: Structural entry moments

DiD treatment effect on streaming entry	+47.3 pp
Additional director-credits attributed to the cap	2,153

Notes. Panel A is a linear probability model on a balanced director-by-year panel, $t \in \{2018, \dots, 2022\}$. (The 2017 vintage of m0202_drama_va.csv has missing director identifiers from an upstream merge and is therefore excluded.) The cohort comprises directors with ≥ 2 directing credits in 02_rdata/20_director_creations recorded for year < 2017. The treated group consists of cohort members who directed no streaming drama in 2018 (the single available pre-cap year); the control group consists of those who directed at least one streaming drama in 2018. The outcome $\text{active}_{i,t}$ equals one if director i directs at least one streaming drama with release year t in 03_wdata/m0202_drama_va.csv, else zero. $\text{Post}_t = \mathbb{1}\{t \geq 2019\}$.

Columns (2)–(4) absorb the time-invariant LockedOut_i level via director fixed effects, so only the interaction is identified. Column (3) restricts the cohort to "established" directors with at least five pre-2017 directing credits. Column (4) weights observations by the number of pre-2017 directing credits. Standard errors clustered at the director level in parentheses. $^{*}, ^{**}, ^{***}$: $p < 0.10, 0.05, 0.01$.

Panel B reports medians by cohort and shows that locked-out status is not random: the outside cohort is more likely to be high-quality but also more likely to come from genre categories outside the 2018 streaming core's modal categories. Panel C separates the incumbent-retreat margin from the outside-cohort absorption margin; the 2018 incumbent activity rate is mechanical because the control cohort is defined by 2018 streaming activity. Some locked-out directors are naturally declining in all-media activity, so the post-cap claim is that streaming absorbs an active subset of the outside cohort, not that the entire cohort enters.

Panel D gives the two director-entry moments used in the structural interpretation.

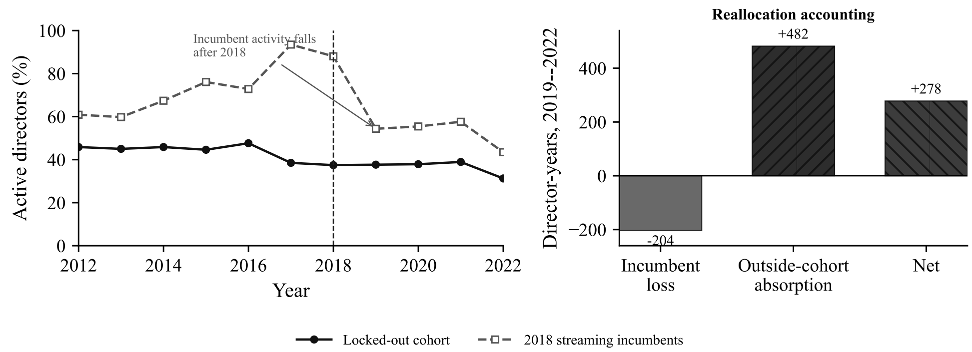

Figure 3: Director reallocation around the cap. The left panel plots each 2018-defined cohort's annual activity in the long-run director filmography, beginning in 2012; this panel is a cohort-identity diagnostic, not a streaming-entry pre-trend test. The right panel reports director-year accounting over 2019–2022: outside-cohort absorption relative to zero baseline, incumbent loss relative to repeating 2018 incumbent activity, and the net. The locked-out cohort consists of pre-credentialed directors with no streaming drama in 2018; incumbents directed at least one streaming drama in 2018.

by international and state-media reporting to inflated contracts and "yin-yang" contracts; contemporaneous coverage notes that Chinese authorities introduced salary ceilings for film and television and blamed sky-high salaries and dual contracts for tax evasion and a money-worship culture (Meixler, 2018; Desta, 2018; Yuan, 2019). The later Zheng Shuang case provides a post-cap enforcement example: state-tax and broadcast-regulator reports described a 160 million RMB agreed compensation package for *A Chinese Ghost Story*, under-declared income, and follow-on broadcast sanctions (State Taxation Administration of China, 2021; National Radio and Television Administration, 2021). Appendix Table 26 reports these sources and the iQIYI advertising-revenue background in one place. These sources establish that the industry and regulators recognized extreme star pay and attention monetization as real institutional facts. What they do not prove is the economic mechanism. That is why Figure 4 is needed.

In our data, high-traffic stars are privately valuable inputs for platforms. The key empirical question is not only whether rewards are concentrated, but whether chasing stars is a rational platform choice. Figure 4 shows the visual pattern, and Table 6 reports the corresponding regression and concentration moments.

Table 6 should be read in three steps. First, Panel A asks whether buying star input is privately rational for platforms before the cap. The positive quadratic coefficient, 0.060, says that the baseline payoff curve is convex: higher star input is associated with more than proportional view payoff. After adding release-year, platform, genre, and length

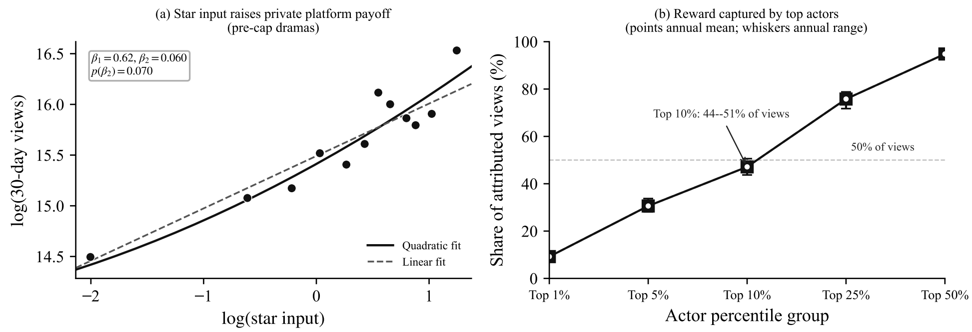

Figure 4: Why platforms bid for traffic stars. Panel (a): pre-cap drama-level relationship between star input and 30-day views; points are binned means and lines are fitted returns. Panel (b): cumulative share of attention-attributed views captured by top actor groups; points are annual means and whiskers are annual ranges.

Table 6: Payoff Convexity and the Star-Tournament Mechanism

Moment / regression test	Estimate	SE
Panel A: Is buying star input privately rational for platforms?		
Pre-cap quadratic: $[\log(\text{star input})]^2$	0.060	0.033
With year, platform, genre, and length controls	0.082	0.024
Extra slope above the pre-cap 90th percentile	1.629	0.904
Top-10 star-input bin premium vs bottom half	1.259	0.165
Panel B: Is the payoff schedule durable rather than a pre-period artifact?		
Post-cap quadratic using frozen pre-2018 star input	0.088	0.045
Panel C: Can stars convert convex payoff into bargaining rents?		
Actor reward rank slope, with year fixed effects	-1.961	0.016
Top-10 actor share of attributed views	47.1%	2.1%
Pre-cap drama observations	615	
Actor-year observations	43,507	

Notes. Panels A–B use the Yunhe–Douban drama panel and regress log(30-day views) on pre-2018 maximum actor view-VA among the top-five billed cast. Controls are release-year, main-platform, and macro-genre fixed effects plus log episode count. The top-spline term is the extra slope above the pre-cap 90th percentile. Panel C uses the actor–drama value-added panel and constructs actor-year attributed views as $\sum_j$ attention share$_{ij}$ × 30-day views$_j$. Robust (HC1) standard errors are reported. The table is a mechanism diagnostic: it shows private returns to star input and concentration of attention rewards, but does not by itself observe wage bids or bargaining offers.

controls, the coefficient rises to 0.082, so the convexity is not just a platform, genre, year, or duration artifact. The top-tail spline then shows that the payoff slope steepens further above the pre-cap 90th percentile. Finally, the top-10 star-input bin premium of 1.259 log points means that these projects receive roughly 3.5× the view payoff of bottom-half projects, conditional on controls. This is a private-payoff test, not a direct observation of bidding: stars may also proxy unobserved promotion budgets, IP quality, or production scale. The role of the test is to show that chasing traffic stars is rational for platforms under the pre-cap attention technology.

Second, Panel B asks whether this payoff schedule is only a pre-cap artifact. The post-cap quadratic, computed using frozen pre-2018 star input, remains positive at 0.088. The interpretation is that audience demand for traffic-star attention did not disappear after the cap. What changes in the model is the bargaining environment: the cap weakens the star's streaming outside option, so platforms can capture more residual surplus even though stars remain valuable demand inputs.

Third, Panel C asks whether stars can convert the convex payoff schedule into bargaining rents. The actor reward rank slope is −1.961, implying a steep winner-take-most distribution of attributed views: moving down the rank distribution is associated with a sharp fall in attention. The top 10% of actors capture 47.1% of attributed views on average, with annual shares ranging from 43.7% to 50.6%. This scarcity is what makes a Bertrand-style outside option plausible. Together, the three panels support the star-tournament mechanism: traffic stars raise expected platform payoff convexly, and the scarcity of those stars makes bargaining rents feasible. The auction interpretation itself is completed structurally by the residual-surplus and make-or-buy moments, which estimate the pre-cap outside-option intensity rather than assuming it from this reduced-form table alone.

A natural concern is that a tournament should reduce the number of series, while some external platform and advertising measures show more projects after the cap. The distinction is between released-title supply and lottery-ticket entry. The main Yunhe–Douban released-title panel does not show a simple quantity boom: tracked titles fall from 466 in 2018 to 400 in 2024. But external validation sources show many more attempts below the blockbuster tier. The Endata sponsorship sample rises from 51 dramas in 2018 to 125 per year on average during 2019–2022, while the long-tail share of low-sponsorship dramas nearly doubles and top-10 brand-exposure share remains concentrated. iQIYI's public production pipeline similarly expands from 32 announcements in 2015–2018 to 367 in 2019–2025. Appendix Table 27 reports these moments. The interpretation is not "more quantity means no tournament." It is that a winner-take-most payoff schedule can generate many attempts to create hits, while realized attention and sponsorship remain concentrated at the top.

After the policy, top stars become easier to afford within surviving projects. Top-star coverage and intensity rise, while newcomer exposure falls. This project-level cast recomposition is a held-out tension moment rather than a targeted fit: the current structural model has one lead-role matching margin, not a full multi-slot cast portfolio. At the same time, some high-reputation actors appear to reallocate away from high-quality lead roles and toward endorsements and other commercial outside options. The endorsement evidence is informative but narrower than the cast panel: Endata covers a matched commercial sample, so endorsement counts should be read as an outside-option proxy, not as population endorsement rates or direct replacement wages. Figure 5 documents the project-level affordability margin, while Table 7 documents the within-actor career-shift margin.

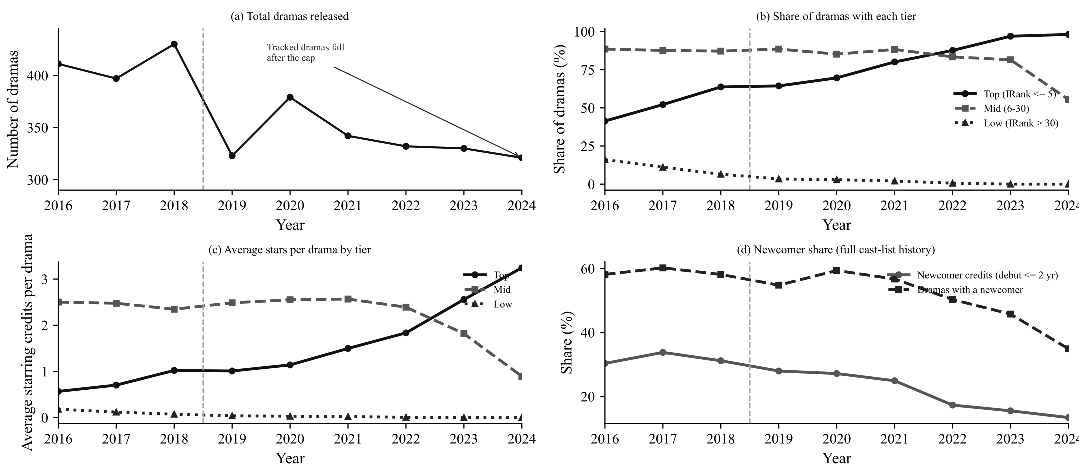

Figure 5: Top-star coverage and actor reallocation. Panel (a): total tracked dramas. Panel (b): share of dramas with at least one top-, mid-, or low-tier actor. Panel (c): average starring credits per drama by tier. Panel (d): newcomer exposure based on full cast-list history.

3.5 Market output, variety, and selection

The final effect is on market output and the margins of variety. Variety is not a single object in this market. The evidence therefore has three layers, each shown explicitly in Figure 6 and Table 8. First, the producer margin broadens. Distinct directors in the released-drama panel rise from 97 in 2018 to a 2019–2022 average of 192, while director HHI falls from 0.012 to about 0.006. This is the reduced-form counterpart of the production-frontier mechanism: the cap is associated with more producer identities reaching the platform market.

Second, category variety does not mechanically expand. Active Douban genres fall from 16 in 2018 to a post-cap average of 14, and genre HHI rises from 0.173 to about 0.193.

Table 7: Within-Actor Career Shift

Pre-rating bin	N	$\Delta P(H)$	$\Delta P(L)$	E_{pre}	E_{post}	ΔE	Substitution corr. $\rho(\Delta E, \Delta H)$
Low	1099	+0.87	-2.61	0.35	1.15	+0.80	+0.161
Mid	1134	+0.09	-1.55	0.15	0.70	+0.55	+0.240
High	1065	-2.75	+2.07	0.06	0.33	+0.26	-0.182

Sample: 3,298 actors active in both 2017–2018 and 2019–2022. Pre-2018 rating bin is a tertile split on the actor's pre-2018 rating value-added. Composition margins are own-actor differences between post- and pre-period shares of drama appearances that are top-2 billed in a high-Douban (≥ 7) or low-Douban (< 7) production.
Endorsement counts are summed brand contracts started in the indicated four-year window (*Pre*: 2015–2018; *Post*: 2019–2022), from Endata.
Key reading: high-rating-reputation actors lose high-quality lead exposure by 2.75pp and gain low-quality lead exposure by 2.07pp, while their endorsement-substitution correlation flips sign ($\rho < 0$) relative to lower-reputation actors ($\rho > 0$). Here H denotes high-quality lead roles, L denotes low-quality lead roles, and E denotes matched endorsement contracts. High-reputation actors who land more matched endorsements take *fewer* high-quality lead roles. Because endorsement coverage is sparse, the endorsement columns should be read as matched-sample outside-option proxies, not population endorsement rates.

Appendix Figure 17 and Tables 22–23 show the composition behind this concentration: the market shifts away from Comedy, Fantasy, Sci-Fi, and Action and toward Crime, Drama, History, Mystery, and Costume Drama, with much of the shift coming from the entry-margin directors documented above. Thus the cap expands *who* produces, not necessarily the number of genre categories produced.

Third, the released market becomes more selected. Total tracked dramas fall from 466 in 2018 to 400 in 2024, while observed Douban quality improves gradually: the mean score is 5.77 in 2018, remains around 5.7–6.0 during 2019–2021, and then rises to 6.82 by 2024. The Douban-score IQR also narrows from 2.225 in 2018 to about 1.863 during 2019–2022. Figure 7 reports the corresponding quantity and score-distribution paths. This is a held-out validation fact about market selection, not a sharp project-level treatment effect. The baseline simulation is asked to match total drama supply, while full consumer-side variety remains a structural welfare object under the demand model rather than a directly observed reduced-form claim.

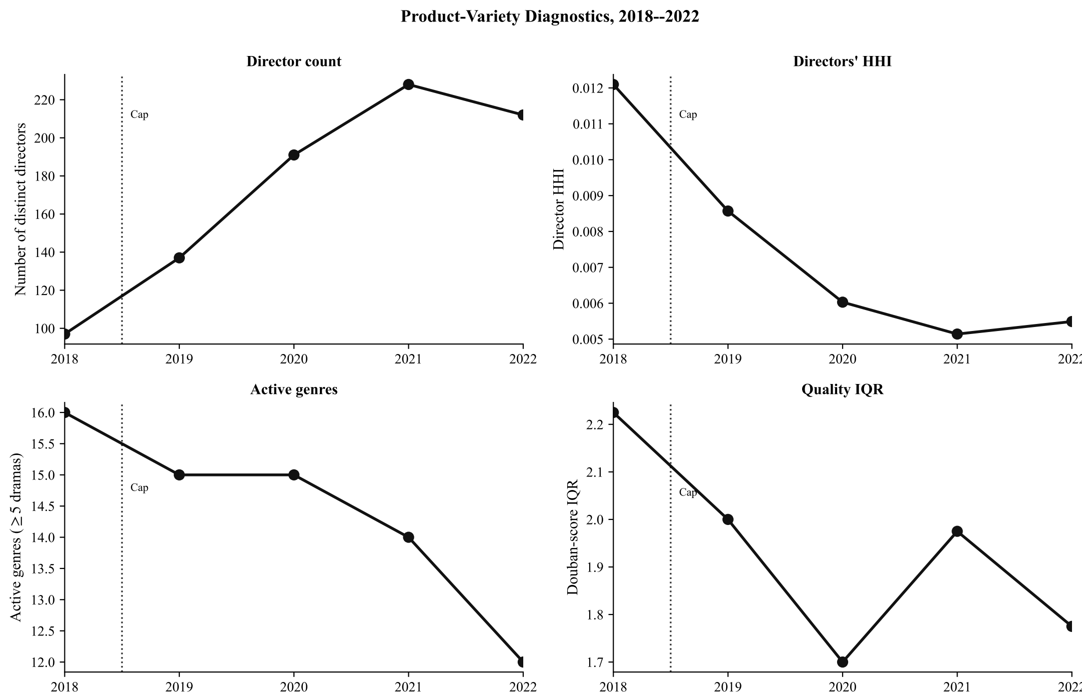

Figure 6: Fact 6, layers of variety. Panel (a) plots distinct directors, the producer-participation margin. Panel (b) plots director HHI, the producer-concentration margin. Panel (c) plots active Douban genres, the content-category margin. Panel (d) plots the Douban-score interquartile range, the quality-selection margin. The vertical dashed line marks the post-cap period.

Table 8: Product-Variety Diagnostics

	2018	2019	2020	2021	2022	Δ(post–2018)
Total dramas	466	421	454	434	410	
Distinct directors	97	137	191	228	212	+95
Director HHI	0.012	0.009	0.006	0.005	0.005	-0.006
Active genres (≥ 5 dramas)	16	15	15	14	12	-2
Genre HHI	0.173	0.187	0.171	0.201	0.212	+0.020
Distinct top-1 leads	338	296	341	331	304	-20
Top-1 lead HHI	0.003	0.004	0.003	0.004	0.004	+0.000
Douban-score IQR	2.225	2.000	1.700	1.975	1.775	-0.363

Notes: All measures computed at the release-year level on the drama panel for 2018–2022. (The 2017 vintage of m0202_drama_va.csv has director identifiers missing for an upstream merge reason and is therefore excluded.) The HHI is the Herfindahl-Hirschman Index on the entity's share of within-year drama production; lower values indicate more diversification. "Active genres" counts Douban genres with at least five dramas in the year. Douban-score IQR is the interquartile range of drama scores in the year. The Δ column reports the post-cap (2019–2022) average minus the pre-cap (2018) value.

Reading: The variety pattern is mixed, not a uniform expansion. The producer margin broadens: distinct directors rise from 97 in 2018 to a 2019–2022 average of 192, and director HHI falls from 0.012 to about 0.006. But genre variety narrows: active genres fall from 16 to a post-cap average of 14, and genre HHI rises from 0.173 to about 0.193. Top-1 lead concentration is approximately flat, while the Douban-score IQR narrows from 2.225 to about 1.863. Thus the reduced-form fact is best read as producer diversification combined with stronger market selection and genre concentration. Full consumer-side variety is therefore a structural welfare object under the BLP demand model, not a directly observed reduced-form fact.

Panel (c) of Figure 7 keeps the overall quality distribution and overlays project-control categories. The classification first uses observed drama-level sourcing labels, then exact-title matched internal acquisition labels, and finally fills remaining missing labels with a deterministic text-based rule aligned with the platform-year project-control shares used in the make-or-buy block. This completion is important because public sourcing fields are sparse in some years, so the category overlay should be read as an auxiliary diagnostic rather than a separately identifying moment.

Reading the evidence as a platform-market mechanism. The empirical facts line up with the broader platform-economy implication in different strengths. The strongest evidence is on the transmission from platform monetization to upstream input rents: advertising and subscriber acquisition make attention privately valuable to platforms; pre-cap star input predicts 30-day views convexly; and actor attention rewards are highly concentrated. The second strong margin is the residual channel: iQIYI's content-cost residual is negative before the cap and recovers afterward, while the structural bargaining block estimates a

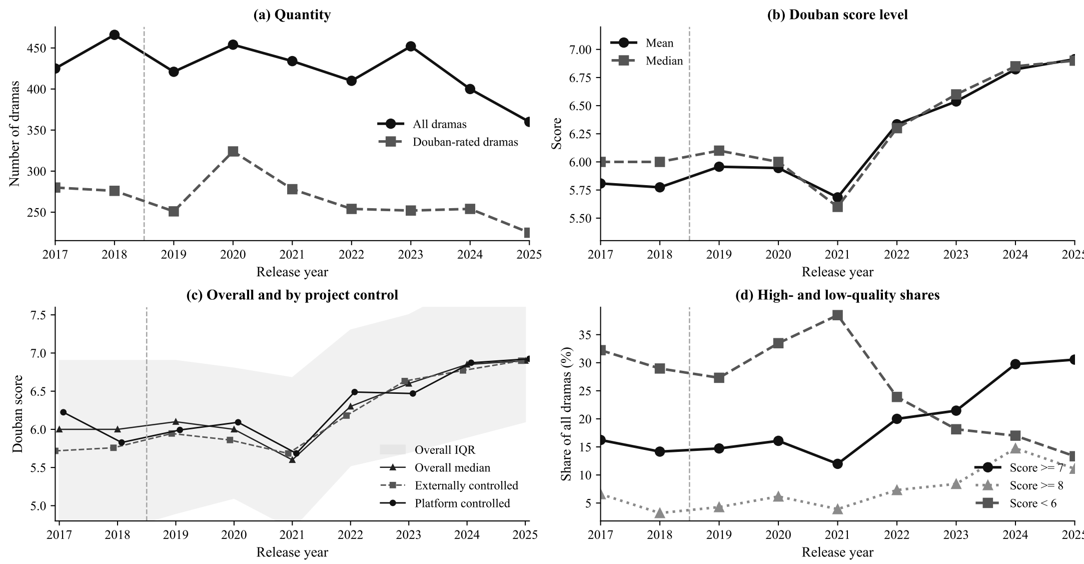

Figure 7: Quantity and quality distribution. Panel (a): total and Douban-rated dramas in the structural Yunhe-Douban panel. Panel (b): mean and median Douban score. Panel (c): overall interquartile range and median score, with project-control category means overlaid. Project control is classified from observed self-produced/licensed sourcing labels, supplemented by internal acquisition labels, and completed by a text-based assignment aligned with platform-year project-control shares. Panel (d): shares of high- and low-scoring dramas. The vertical dashed line marks the September 2018 cap announcement.

near-Bertrand pre-cap outside option. The third strong margin is production organization: audited content-cost flows and three-platform project-control shares show a shift toward platform-controlled production after the cap. The director-entry, pipeline, and variety diagnostics support the production-frontier interpretation: the post-cap platform pipeline expands, more producer identities reach the released market, and many entrants come from outside the pre-cap streaming core. The content-variety claim is more nuanced. Released-title counts fall, observed quality rises, and genre concentration increases, while sponsorship and pipeline data show more long-tail attempts below the blockbuster tier. We therefore interpret the evidence as rejecting a pure wage-transfer story and supporting a platform-induced upstream production distortion, while treating full consumer-variety welfare as a counterfactual object rather than a directly observed reduced-form fact.

4 Model and Identification

This section asks what minimal model can rationalize the joint empirical pattern and, block by block, states which moments discipline each object. If the cap only lowered the price of actors, a static wage model would be enough. But the data move on margins that require equilibrium objects: star marginal revenue product, bargaining outside options, platform control rights, director entry costs, actor career dynamics, and matching. The model therefore targets the margins needed for counterfactuals rather than every descriptive fact in Section 3. For each block below, the model object and the identifying variation are presented together; Section 5 then reports estimates, tuning choices, and fit rather than restating the model.

The model has five agents—consumers, platforms, directors, actors, and advertisers—in a dynamic game. Each subsection states the behavioral object, the empirical moment that identifies or disciplines it, and how the object feeds the next layer. Formal payoff matrices and dominance proofs are in Appendix A.

Industry quantity and selection. The model separates the salary-cap mechanism from a contemporaneous industry extensive margin. The pre-cap market contained many low-cost experimental dramas, and these projects often did not rely on top stars or high-reputation actors. Thus a higher pre-cap count of new series is not evidence that star-intensive production was cheaper or more efficient before the cap. We represent this background dynamic with a yearly external-supply shifter μ_t that governs the mass of candidate externally controlled projects:

$$J^0_{\mathrm{buy},t} = \exp(\mu_t), \qquad J_{\mathrm{buy},t} = J^0_{\mathrm{buy},t}\ \Pr(V_k - c_k \geq \chi_t)\,, \tag{1}$$

where $V_k - c_k$ is the external studio's quality-cost wedge and χ_t is the market selection threshold. In the baseline estimation, μ_t is treated as an industry-time shifter rather than as a salary-cap primitive: demand includes release-year controls, and the supply simulation uses the observed external-drama path when estimating the bargaining and project-control primitives. This choice prevents the model from mechanically attributing the post-cap decline in total new series to the cap. The cap's direct structural channel is instead the bargaining-regime switch that changes wages, platform residual surplus, control rights, and director entry. The quantity and quality facts then serve as validation of the selection margin: the industry moves from broad, low-cost experimentation toward fewer, more selected projects.

4.1 Identification model

This subsection states the minimal empirical model that makes the salary-cap episode informative before adding the full structural machinery. Let C_t^R denote the release-year post-cap regime used for aggregate platform-year and market-year moments. Let C_j^P denote project-production exposure, equal to one when public sources place the filming start after the September 2018 cutoff, zero for projects completed before the cutoff, and missing or partial for projects with unavailable or straddling production windows. The timing audit in Table 1 shows why the distinction matters: release-year post status overstates clean project-level exposure for some 2019 releases, so C_t^R should be read as a regime indicator and C_j^P as the cleaner project-level treatment when production timing is available. A platform-project-year outcome Y_{jpt} can be written as

$$Y_{jpt} = \beta_C C_{jpt} + \beta_A A_t + X'_{jpt}\beta_X + \alpha_p + \tau_t + \varepsilon_{jpt}, \tag{2}$$

where C_{jpt} is C_t^R for aggregate moments and C_j^P for project-level timing tests, A_t is the contemporaneous long-form attention shock, X_{jpt} contains drama, genre, and platform controls, α_p are platform fixed effects, and τ_t captures common industry-time pressure. Equation (2) is not the estimating equation for every block. It is the reduced-form object that clarifies what the cap must explain: a change in the bargaining environment should move a linked set of margins in the same direction, while a generic attention shock or industry contraction need not.

The commitment channel has four testable implications. First, if the cap weakens rival-platform bidding for scarce stars, platform residual surplus should recover after the policy:

$$\gamma_{pt} = \gamma_0 + \rho C_t^R + \psi A_t + \alpha_p + u_{pt}, \qquad \rho > 0. \tag{3}$$

Second, if the recovered residual finances control rights, the platform-controlled project share should rise:

$$\Pr(M_{jpt} = 1) = \Lambda\Big(\delta_0 + \delta_\gamma \gamma_{pt} + \delta_C C_{jpt} + \delta_A A_t + X'_{jpt}\delta_X + \alpha_p\Big), \qquad \delta_\gamma > 0. \tag{4}$$

Third, if platform control expands the feasible production frontier, previously locked-out experienced directors should enter streaming relative to 2018 incumbents:

$$Enter_{dt} = \theta_d + \theta_t + \pi\left(LockedOut_d \times Post_t\right) + e_{dt}, \qquad \pi > 0. \tag{5}$$

Fourth, if the transition is not merely lower star pay, the market should show selection in organization and output: fewer low-residual external projects, higher platform-control intensity, and a rightward shift in observed quality among surviving projects.

These equations define the main threats. A short-form attention shock can explain a contraction in long-form output, but by itself predicts lower revenue pressure rather than a recovery of platform residual surplus and a rise in platform-controlled production. A subscription pivot can change revenue composition, but it does not directly predict the entry of previously locked-out directors following a failed private salary agreement. A generic regulatory tightening can reduce production volume, but it does not uniquely predict the linked residual-control-entry pattern. The structural model below therefore does not replace identification. It quantifies the no-cap counterfactual conditional on the reduced-form fact that the policy moved the bargaining and control-rights margins together.

4.2 Consumer demand

Consumers choose at most one drama per market m (release year); they may also pick the outside option $j = 0$. Consumer i's indirect utility from drama j is

$$u_{ijm} = X_j\beta_i + \alpha_i p_{jm} + \xi_j + \epsilon_{ijm}, \quad \begin{pmatrix}\alpha_i \\ \beta_i\end{pmatrix} = \begin{pmatrix}\alpha \\ \beta\end{pmatrix} + \Sigma v_i, \quad v_i \sim N(0, I), \tag{6}$$

where X_j collects drama characteristics (genre, episode count, director value-added, and cast reputation); ξ_j is unobserved drama quality; ϵ_{ijm} is i.i.d. Type-I extreme value; and (α_i, β_i) are random taste coefficients with population means (α, β) and dispersion Σ.

The price p_{jm} is the platform's monthly VIP subscription rate in market m, varying primarily at the platform×year level. The empirical specification absorbs p_{jm} into platform×year fixed effects and instruments the platform-level subscription price using a Hausman instrument equal to the mean subscription price of competing platforms in the same year. Integrating over v_i yields shares s_j and—via $\partial s_j / \partial R_a^P$—each actor's marginal revenue contribution V_a, the single demand-side input to the platform–talent bargaining block.

4.3 Producer

A potential project is organized around a director d with type $\tau_d \in \{\text{idol}, \text{narrative}\}$ (idol drama versus narrative—crime, mystery, history). The platform–director production unit makes two sequential decisions: entry, then cast portfolio.

Entry is a binary $\text{Enter}_{d,t} \in \{0, 1\}$: director d enters streaming in year t if the platform-financed streaming payoff exceeds the director's outside option $\Omega_{d,t}$ from adjacent media:

$$\pi_{d,t}^{\text{entry}} \;=\; \mathbb{E}\big[\, \gamma_{p,t} \cdot S(A_j^*, \tau_d, \bar{W}_t) \,\big] \;-\; \kappa_d, \tag{7}$$

where $S(A_j^*, \tau_d, \bar{W}_t) = R(s_j) - C_{\text{prod}}(\tau_d) - \sum_a W_a^*$ is project surplus under the optimal cast A_j^*; $\gamma_{p,t}$ is the platform residual share available to finance project control and entry; and κ_d is the director's streaming-specific fixed entry cost. With $z_{d,t} \;=\; \pi_{d,t}^{\text{entry}} - \Omega_{d,t}$ and $\varepsilon_{d,t} \sim N(0, \sigma_\varepsilon^2)$,

$$\Pr(\text{Enter}_{d,t} = 1 \mid \mathcal{F}_{d,t}) \;=\; \Phi\!\left(\frac{z_{d,t}}{\sigma_\varepsilon}\right), \tag{8}$$

estimated by MLE in Section 5.5. The cap shifts $z_{d,t}$ through two channels: $\gamma_{p,t}$ rises from 0 (Bertrand) to $\omega(1 - d_S/V_a)$ (Nash), raising the platform residual share; and the residual non-cast budget $\bar{B} - W^*$ expands until it covers the narrative non-cast cost $C_{\text{prod}}(\text{narrative}) = \bar{c}$, an incentive-compatibility flip that admits narrative-type directors who optimally stayed out pre-cap.

Conditional on entry, the platform–director unit picks cast A_j to maximize residual surplus subject to the two-tiered cap:

$$\begin{aligned} &\max_{A_j} \;\; \gamma_{p,j} \left[R\big(s_j(A_j)\big) - C_{\text{prod}}(\tau_d) - \textstyle\sum_{a \in A_j} W_a^* \right] \\ &\text{s.t.} \quad \textstyle\sum_{a \in A_j} W_a^* \le \bar{W}_{\text{total}} \text{ (C1)}, \quad W_a^* \le \bar{W}_{\text{ind}} \;\; \forall a \text{ (C2)}. \end{aligned} \tag{9}$$

At $\widehat{\omega} = 0.599$ the Nash wage implies a post-cap actor revenue share of $W/R = 0.234$. In the production-cost denominator used by the regulation, the disclosed iQIYI projects have a median lead-actor share of 17.2% and a maximum of 25.3%, below the 40% total-actor-pay ceiling. The cap therefore primarily reshapes cast portfolios by switching the upstream bargaining regime, rather than by continuously binding inside the Nash regime.

The star's bargaining position is therefore partly endogenous and partly a maintained bargaining primitive. The endogenous part is the actor's realized surplus claim: it depends on the BLP-implied marginal revenue product V_a, the actor's outside option d_S, the pre-cap Bertrand outside-option intensity φ, and the post-cap opportunity set generated by the lead-arrival and matching layers. High-V_a actors with scarce substitutes have stronger realized bargaining positions because rival platforms value the same attention asset. The maintained primitive is the Nash weight ω inside the post-cap bilateral bargain. We do not derive ω from a separate microfoundation of negotiation protocols, agency contracts, or platform liquidity; instead, we calibrate it from the disclosed post-cap wage-share moment, conditional on V_a and d_S. Thus the paper should be read as endogenizing the economic objects that make stars valuable and scarce, while treating the post-cap Nash split parameter as a transparent calibrated primitive.

Proposition 1. *A binding salary cap that lowers the star wage from the Bertrand level V_a to the Nash level $W^* < V_a$ expands the set of feasible director types from idol-only to idol-plus-narrative and lowers the equilibrium industry HHI.*

4.4 Actor career path

Actors are forward-looking, but they do not freely choose every role tier in every period. They first face an opportunity process, then decide whether to accept the available role. An actor in state $S_{at} = (R^Q_{at}, R^P_{at})$ (quality- and popularity-reputation stocks) receives a lead-role opportunity with probability

$$\begin{aligned}\lambda_{at} = \Lambda\big(&\eta_0 + \eta_Q R^Q_{at} + \eta_P R^P_{at} + \eta_F \text{Flow}_{at} + \eta_B \text{Brand}_{at} \\ &+ \eta_N \log(1 + N^{\text{prior}}_{at}) + \eta_G \text{Brokerage}_{at} + \eta_X X^{\text{match}}_{at}\big),\end{aligned} \tag{10}$$

where $\Lambda(\cdot)$ is the logistic CDF. This is the model's extensive-margin object. It summarizes market access, visibility, availability, agency intermediation, and search frictions: high reputation and popularity affect whether the actor is considered; flow and brand exposure proxy current visibility; prior dramas proxy career embeddedness; brokerage captures access to large casting networks; and X^{match}_{at} is the matching-implied exposure index from the director-slot market.

Conditional on a lead opportunity arriving, the director–actor matching layer determines whether the opportunity is high- or low-quality:

$$q_{at} \equiv \Pr(\text{lead\_high}_{at} = 1 \mid \text{lead arrival}, S_{at}) = \sum_{d \in \mathcal{D}^H_t} \frac{\exp\{\Phi(x_{at}, y_{dt})\}}{\sum_{d' \in \mathcal{D}_t} \exp\{\Phi(x_{at}, y_{d't})\}}, \tag{11}$$

with $\Phi(x, y) = x^\top A y$. Thus the extensive margin λ_{at} governs whether the actor receives a lead opportunity at all, while q_{at} governs the composition of that opportunity. Supporting opportunities are treated separately because the current matching block is a lead-role matching model.

Let $C(O_{at})$ denote the feasible action set implied by the realized opportunity $O_{at} \in \{\text{none}, \text{supporting}, \text{lead\_low}, \text{lead\_high}\}$. The actor chooses among feasible actions by solving

$$V(S_{at}) = \mathbb{E}_{O_{at}}\left[\max_{\text{Role}_t \in C(O_{at})}\left\{W^\star_{at}(\text{Role}_t, S_{at}) + E_{at}(R^Q_{at}, R^P_{at}) + \beta\,\mathbb{E}[V(S_{a,t+1}) \mid S_{at}, \text{Role}_t]\right\}\right], \tag{12}$$

with $R^k_{a,t+1} = (1-\delta_k)R^k_{at} + f_k(\text{Role}_t) + \eta^k_{at}$ for $k \in \{Q, P\}$. The actor trades current income against reputation investment: high-visibility leads build R^P, narrative leads build R^Q.

The flow wage $W^\star_{at}$ is the equilibrium output of the bargaining block, evaluated at the actor's state and role:

$$W^\star_{at} = \begin{cases} V_a(S_{at}, \text{Role}_t) & \text{(pre-cap Bertrand)}, \\ (1-\omega)V_a(S_{at}, \text{Role}_t) + \omega d_S & \text{(post-cap Nash)}. \end{cases} \tag{13}$$

Endorsement $E_{at}(R^Q, R^P)$ is a passive flow in both reputation stocks; we do not model an endorsement capacity constraint, a limitation we return to in Section 8.

4.5 Platform and talent: Nash bargaining

Each platform commissions a portfolio in year t with profit

$$\Pi_{kt} = \sum_{j \in \mathcal{J}_{kt}}\Big[\big(\lambda_t ARPU_{adv,t} + (1-\lambda_t)ARPU_{sub,t}\big)M_t\, s_{jt} - C_{jt}\Big] - F_{kt}, \tag{14}$$

where λ_t is the ad-versus-subscription revenue mix, M_t is market size, s_{jt} is drama share, and F_{kt} is overhead. The cap touches only the wage bill.

Both bargaining regimes are anchored to demand-derived value-added:

$$V_a \;\equiv\; \mathbb{E}_m\left[\big(\lambda_t ARPU_{adv,t} + (1-\lambda_t)ARPU_{sub,t}\big)M_t \cdot \Delta s_{jm}(R^P_a)\right], \tag{15}$$

where $\Delta s_{jm}(R^P_a)$ is the BLP share elasticity to swapping actor a for the next-best alternative. V_a is the platform's reservation wage; what is actually paid depends on the bargaining mechanism.

Pre-cap, the star's high realized bargaining power is generated by the outside option rather than by assigning the actor a larger Nash weight. Let d^0_S denote the actor's non-streaming fallback and let $\varphi \in [0,1]$ summarize how strongly rival-platform bidding capitalizes the star's attention value into the actor's fallback:

$$d^{\text{eff,pre}}_S \;=\; \varphi V_a + (1-\varphi)d^0_S. \tag{16}$$

This object is endogenous to the market environment in the following sense. V_a comes from consumer demand and platform monetization; scarcity enters because several platforms want the same attention asset; and φ is estimated from the pre-cap platform-residual moment together with the three-platform project-control transition, rather than chosen directly to fit wages. When φ is close to one, rejecting platform p is almost equivalent to selling the same attention asset to a rival platform. The actor therefore has a high effective disagreement payoff even if the primitive Nash weight ω is unchanged.

In the Bertrand limit with $K \geq 2$ platforms able to bid for the same scarce star, standard undercutting drives the equilibrium wage to its ceiling,

$$W_a^{\text{Bertrand}} \;=\; V_a, \tag{17}$$

so the platform's share is $\gamma_p^{\text{Bertrand}} = 0$. Top stars empirically absorbed 50–70% of production cost; iQIYI's accounting residual $\widehat{\gamma}_{p,2018} = -8.6\%$ reflects fixed and leverage costs on top of the bilateral $\gamma_p = 0$.

The cap $W \leq \bar{W}_{\text{ind}} < V_a$ breaks the auction's enforcement chain—the rival cannot pay V_a either—collapsing the regime to bilateral Nash bargaining. The same switch realizes in production mode through control rights. Pre-cap, platforms rely more heavily on content whose production organization and talent contracting are controlled by external studios; post-cap, they move toward platform-controlled projects in which the platform either produces through an affiliate or commissions an external studio under platform control. Thus "make" below should be read as a control-rights state, not as literal physical production by platform employees.

Platform bargaining weight $\omega \in (0, 1)$; disagreement payoffs are $d_P = 0$ and d_S. The star outside option d_S is the actor's best fallback if she does not accept the current platform's lead-role contract. Economically, it includes non-streaming acting work, film, variety, endorsements, commercial appearances, and the continuation value of waiting for a better career path. In the baseline bargaining block we do not separately model each of these markets; $d_S/R = 0.040$ is a reduced-form post-cap fallback calibrated as a share of project revenue. Pre-cap, rival streaming platforms add the competitive component in Equation (16), making the effective outside option much larger. The Nash problem is

$$W_a^* \;=\; \arg\max_W \;\left[(V_a - W) - d_P\right]^{\omega}\left[W - d_S\right]^{1-\omega}, \tag{18}$$

with first-order condition

$$W_a^* \;=\; (1-\omega)V_a + \omega d_S, \qquad \gamma_p^* \;=\; \omega\left(1 - \frac{d_S}{V_a}\right). \tag{19}$$

We treat ω as a post-cap bargaining primitive calibrated from the disclosed wage-share moment; pre-cap data cannot identify it because the Bertrand auction makes the Nash split irrelevant. Within Nash, the central iQIYI revenue-share moment gives $\widehat{\omega} = 0.599$ and $W^*/R_j = 0.234$; in the regulation's production-cost denominator, disclosed lead-actor shares remain below the 40% total-actor-pay ceiling. The cap binds primarily as a regime-switching mechanism.

Figure 8 summarizes the structural accounting. Stars are a scarce core input: the BLP-implied marginal revenue product of one lead star is $V_a/R_j = 0.523$. Because platforms bid against one another in a Bertrand auction, the pre-cap wage is bid up to this level ($W^{\text{Bertrand}}/R_j = V_a/R_j$), leaving the platform no residual surplus. The cap removes the rival's ability to pay V_a, collapsing the regime to bilateral Nash bargaining with $\widehat{\omega} = 0.599$ and calibrated outside option $d_S/R_j = 0.04$. The post-cap Nash wage is therefore $W^*/R_j = (1-\omega)V_a/R_j + \omega d_S/R_j = 0.234$, and the platform captures roughly 55% of the star MRP. The high pre-cap salary is thus endogenous to the star's bargaining position and scarcity, not to an exogenous cost shock; the cap restores a positive platform share by changing the bargaining regime.

Figure 8: Star MRP and Bargaining Split

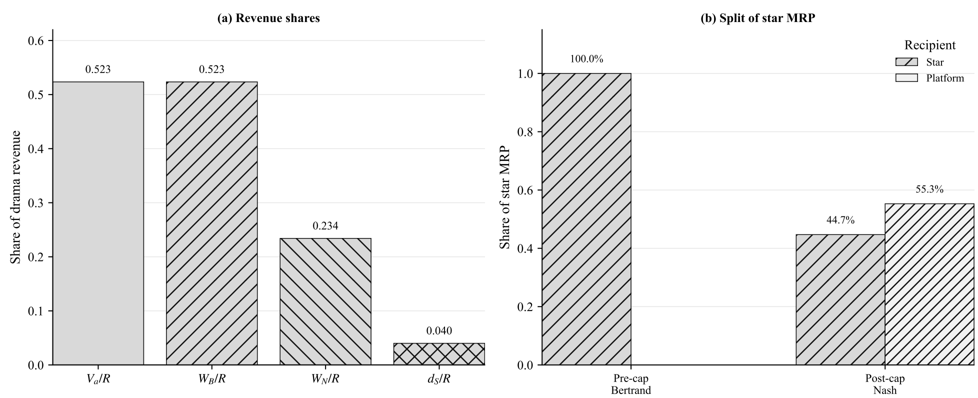

Notes. Panel (a) reports all quantities as shares of drama revenue. V_a/R is the BLP-implied marginal revenue product of one lead star; W_B/R is the pre-cap Bertrand wage, equal to V_a/R when platforms bid for the same scarce star; W_N/R is the post-cap Nash wage implied by $\widehat{\omega} = 0.599$; and d_S/R is the calibrated post-cap actor fallback value. Panel (b) normalizes the star MRP to one and shows who captures it: under Bertrand the star captures the full MRP, while under Nash the platform captures the residual share above the star's wage.

The high MRP is not driven by a handful of superstars. Across post-cap dramas, the median star has $V_a/R_j = 0.523$, the 90th percentile reaches 1.33, and essentially every drama in the sample has V_a/R_j above both the post-cap Nash wage share (0.234) and the non-streaming outside option (0.040). Stars are therefore a scarce core input with high marginal value throughout the distribution, which is why their pre-cap wages were bid up to MRP and why the cap's regime switch has economy-wide bite. Appendix Figure 21

Table 9: Star Marginal Revenue Product and the Bargaining Split

Quantity	Value
Star MRP (V_a/R)	0.523
Pre-cap Bertrand wage ($W/R \approx V_a/R$)	0.523
Post-cap Nash wage (W/R)	0.234
Outside option (d_S/R)	0.040
Nash weight (ω)	0.599
Surplus split (actor share of MRP)	
Pre-cap Bertrand	1.000
Post-cap Nash	0.447
Platform share post-cap	0.553

Notes. V_a/R is the BLP-implied marginal revenue product of one lead star. Because stars are a scarce core input, V_a/R is large. Pre-cap Bertrand competition bids the wage up to V_a, giving the star 100% of the marginal surplus. The cap restores Nash bargaining with $\omega \approx 0.60$ under the central wage-share moment, so the star keeps 44.7% of MRP and the platform captures the remainder.

and Appendix Table 24 report the full distribution.

Panel B of Table 9 makes the regime switch explicit. Pre-cap, the effective outside option is $d_S^{\text{eff}}/R = \varphi V_a/R + (1-\varphi)d_S/R = 0.523$ because stars can still play platforms against one another with Bertrand intensity $\widehat{\varphi} = 0.998$. The platform's residual share is therefore close to zero, and the per-lead wage is $0.523R$. Post-cap, the cap removes the rival platform as a credible bidder, so d_S^{eff}/R collapses to the non-streaming outside option 0.040, γ_M rises to 0.553, and the per-lead wage falls to $0.234R$. The cap is not a conventional price ceiling; it is a public commitment that changes the bargaining regime.

4.6 Endogenous control-rights regime

The platform chooses the control intensity of each project. Let $M_j = 1$ denote a platform-controlled project: pure self-production and commissioned/custom production both enter this state because the platform holds project control, controls the budget, and contracts directly or through a designated contractor. Let $M_j = 0$ denote externally controlled sourcing: licensed, bought, or jointly procured content whose production organization is primarily set by the outside studio before the platform acquires exhibition or distribution rights. Revenue-sharing and undeveloped-IP observations are excluded from the core binary moment and retained as auxiliary categories.

This interpretation matches the project data. In the iQIYI internal project panel, almost all pure self-produced projects list iQIYI or an affiliate as the producing party, but commissioned projects often list both iQIYI and an external studio. The important distinction is therefore not whether an external firm appears in the involved-party field. It is whether the platform has project-level control over the production process and talent contracting. The structural control share is the share of projects in this platform-controlled state.

The choice depends on the platform residual share $\gamma_{M,pt}$, the platform-specific control-rights cost wedge $\Delta\kappa_p(\Psi_{pt})$, and the capability accumulated through past platform-controlled production. We parameterize the platform-control probability as

$$\begin{aligned} P(M_{jpt}) &= \sigma\left(\gamma_{M,pt} V_a - \Delta\kappa_p(\Psi_{pt}) + \beta \Delta EV_{p,t+1}\right), \\ \Delta\kappa_p(\Psi_{pt}) &= \Delta\kappa_{0p} \exp(-a_\kappa \Psi_{pt}), \\ \Psi_{p,t+1} &= (1-\lambda)\Psi_{pt} + M_{pt}. \end{aligned} \tag{20}$$

where $\gamma_{M,pt} = \omega(1 - d_{S,pt}/V_a)$ is the platform's Nash take, $\Delta\kappa_{0p}$ is platform p's baseline cost of exerting project control relative to external sourcing, a_κ measures how accumulated project-control capability lowers that cost, and λ governs capability depreciation. Pre-cap, the outside option is a convex combination $d_S^{\text{pre}} = \varphi V_a + (1-\varphi) d_S^{\text{base}}$, where φ measures Bertrand intensity. At $\varphi = 1$, the auction pins $W^* = V_a$ and $\gamma_M = 0$, sharply reducing the residual budget available for platform-controlled production. The cap severs the Bertrand fallback, restores $\gamma_M > 0$, frees the residual non-cast budget, and lets accumulated control capability lower $\Delta\kappa_p(\Psi)$ over time. Richer variants can add a cost-pass-through term, but the central aggregate make-share specification does not use or interpret such a parameter.

4.7 Moment map, transparency, and implementation

The structural model is estimated as a sequential pipeline. Within each year, decisions unfold in the following order: the policy regime determines whether platform–talent bargaining is pre-cap Bertrand or post-cap Nash; each candidate director decides whether to enter streaming; the platform chooses whether projects are platform-controlled or externally controlled; the matching layer generates lead-role opportunities and their lead-high/lead-low composition; and actors accept or decline opportunities according to the Bellman value of the induced career path. Stages repeat in the next year with updated reputation stocks and policy variables.

The estimation pipeline is: (i) demand (BLP) $\rightarrow$ $(\alpha, \beta, \Sigma, \xi_j)$ and V_a; (ii) attention outside value $\rightarrow$ the time path of long-form attention; (iii) platform Nash bargaining $\rightarrow$ ω; (iv) endogenous project control $\rightarrow$ φ, a_κ, λ, $\Delta\kappa_{0p}$; (v) director entry $\rightarrow$ κ_d distribution; (vi) structured lead-arrival and matching $\rightarrow$ lead opportunities and actor–director sorting; and (vii) actor Bellman $\rightarrow$ acceptance, career values, and reputation transitions.

The moment map follows the same triangular structure. Demand is disciplined by drama-level market shares and subscription-price variation, which identify the BLP price response and the value of star and director inputs. The attention block uses the fall in long-form video attention share to construct an outside-value path for short-form competition. The bargaining block then uses disclosed post-cap wage shares to calibrate the Nash bargaining weight, taking the BLP-implied star marginal revenue product and the actor outside option as inputs.

The producer-side moments identify the transition from star bidding to platform control. The pre-cap residual moment disciplines the Bertrand-intensity parameter, while the iQIYI, Tencent, and Youku platform-year control-share paths discipline each platform's baseline make-or-buy cost wedge, the capability cost-slope, and the capability-decay rate. Director entry is disciplined by the post-cap entry probit and validated against the locked-out director difference-in-differences estimate. Industry supply and external-buy volume are not primary estimating moments for these primitives; they are reported as validation moments for the downstream supply simulation.

The actor-side moments separate opportunity arrival from conditional sorting and acceptance. Post-cap role shares and the 9-cell reputation transition matrix discipline the actor Bellman and stochastic law of motion. The structured lead-arrival layer is disciplined by accepted lead rates across actor states, while the conditional matching model is disciplined by director-type by actor-state match-cell shares. Endorsements and cast composition are reported as validation or tension moments: they help interpret the outside-option and multi-slot cast margins, but the current baseline does not target them as a full cast-market equilibrium.

The structural estimation is therefore overidentified. We do not count the upstream BLP product-share moments in this comparison: the demand block feeds $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\Sigma}$, ξ_j, and V_a/R into the structural model, rather than entering the SMM moment count. Excluding BLP, the benchmark uses eight targeted or calibrated structural moment blocks. These correspond to 483 reported structural cells—the attention path, bargaining wage share, three-platform make-or-buy paths, actor role shares, actor transition matrix, lead-arrival rates, matching cells, and director-entry moments—or 464 independent cells after probability adding-up constraints. The baseline structural model estimates 60 non-BLP parameters, excluding product-level demand unobservables, validation-only extensions, and numerical tuning constants. The number of structural moments is thus much larger than the number of structural parameters.

Throughout the paper, *targeted* evidence enters estimation directly, *calibrated* evidence pins down one parameter by construction, *validation* evidence evaluates a model implication out of target, and *external/tension* evidence marks a pattern that requires a richer extension. Appendix Table 19 gives the full moment-to-parameter map.

The model should be read in three layers. The first layer, *value creation*, estimates consumer demand and translates star and director inputs into revenue-side value. The second layer, *surplus division and platform control*, maps that value into wages, platform residual surplus, and the project-control decision. The third layer, *career allocation*, determines which actors receive lead opportunities, what type of project those opportunities correspond to, whether actors accept them, and how choices update reputation.

The transparency distinction is as follows. Estimated objects are disciplined by the moments described above and summarized in Appendix Table 19. Tuning and calibration choices are fixed before the main counterfactuals and should be read as measurement or normalization choices. Exogenous inputs shift the environment but are not caused by the salary cap inside the model. Maintained assumptions define the boundary of the current structural exercise.

Value creation. The demand block estimates audience value for drama characteristics and maps star and director inputs into V_a/R. It is estimated from drama market shares using the Hausman price instrument and BLP-style competitor instruments. The attention outside value is different: it captures the time-varying price of long-form attention as short-form video expands. In the baseline counterfactual this attention path is treated as an exogenous shock; later sections decompose it with a macro attention regression.

Surplus division and platform control. The Nash bargaining block maps star marginal revenue product into actor wages and platform residual surplus. The bargaining weight ω is calibrated from the disclosed post-cap wage-share moment, while $d_S/R = 0.040$ is an outside-option calibration. This outside option is not the value of leaving acting; it is the actor's best alternative to the current lead-role contract, combining other projects, adjacent entertainment work, endorsements, commercial exposure, and the continuation value of waiting. Pre-cap Bertrand intensity φ measures how strongly platform competition bids this effective outside option toward V_a; it is estimated from the pre-cap residual moment jointly with platform-control share paths and platform-specific make-buy cost wedges. The make-or-buy block then interprets platform-controlled production as more attractive when residual surplus and capability rise, and estimates the dynamic project-control path by SMM/NFXP on 2017–2024 iQIYI, Tencent, and Youku control shares using 320 random starting points.

Career allocation. The actor side separates extensive and intensive margins. The lead-offer extensive margin is the probability that an actor receives any lead opportunity, estimated as a structured arrival-rate logit using reputation, popularity, visibility/access, career history, and matching exposure. The sorting intensive margin governs conditional lead-high versus lead-low composition and actor–director pairing once a lead opportunity exists; it is estimated by conditional logit on observed matched actor–director pairs. The matching layer therefore governs composition, not the arrival rate. The actor Bellman then maps structured lead-arrival rates and matching-generated high/low composition into acceptance, role choice, lifetime values, and reputation transitions. Supporting roles are outside the lead-role matching model and remain disciplined by the observed supporting margin.

Entry and supply. The director-entry block asks whether restored producer residual surplus covers streaming-specific entry costs. It is estimated by a post-cap probit and validated against the locked-out director difference-in-differences evidence. This block links the bargaining and project-control results to the entry margin of directors who were previously active outside the streaming core.

Exogenous inputs. The baseline treats the salary-cap timing, the short-form attention shock, platform subscription-price variation used in the demand IV, the observed external-buy quantity path in the main simulation, and the pre-existing actor and director reputation stocks as exogenous to the model. These objects shift the environment faced by platforms and actors but are not themselves chosen inside the structural model.

Tuning and calibration choices. Table 10 lists the active tuning and calibration constants in the central run. These choices are not interpreted as causal estimates; they normalize payoffs, stabilize numerical integration, or encode industry anchors.

Table 10: Tuning and Calibration Parameters in the Central Run

Block	Parameter	Value	Interpretation
Demand	Market size per year	5.0×10^{11} impressions	Normalization for BLP market shares.
Demand	Random-coefficient draws	200	Simulation draws for BLP integration.
Demand	BLP random seed	20260603	Reproducibility for simulated taste draws.
Bargaining	Outside option share	$d_S/R = 0.040$	Non-streaming actor fallback used to calibrate ω.
Make-or-buy	Taste-shock scale	$\sigma_{\text{reg}} = 0.30$	Scale of the project-control logit.
Make-or-buy	Cap timing in cost decay	2018.5	Mid-year cap timing used in $\Delta\kappa(t)$.
Make-or-buy	Capability grid	$\Psi \in [0, 10]$, 120 points	Grid for backward induction in the dynamic make-or-buy model.
Make-or-buy	Discount factor	$\beta = 0.95$	Future capability value in the dynamic make-or-buy model.
Make-or-buy	Benchmark starts	320	Random starting points for the three-platform NFXP/SMM objective.
Make-or-buy	Start-draw seed	99	Reproducibility for make-or-buy starting points.
Actor Bellman	Discount factor	$\beta = 0.95$	Actor continuation-value discounting.
Actor Bellman	Logit shock scale	$\sigma_\epsilon = 5.0$ RMB million	Scale of role-choice taste shocks.
Actor Bellman	Outside-option grid $d_S(s)$	$0.4, 1.5, 3.0, 0.8, 3.0, 5.0, 1.5, 5.0, 12.0$	Per-state annual outside option for $(L_L, L_M, L_H, M_L, M_M, M_H, H_L, H_M, H_H)$.
Actor Bellman	Wage anchors	lead-high 35, lead-low 8, supporting 5, idle $d_S(s)$	RMB million flow-payoff anchors before state scaling.
Actor law of motion	Role increments $f_Q(a)$	idle 0, supporting 0.05, lead-low 0.10, lead-high 0.30	Quality-reputation increment by role.
Actor law of motion	Role increments $f_P(a)$	idle 0, supporting 0.10, lead-low 0.20, lead-high 0.50	Popularity-reputation increment by role.
Actor law of motion	Persistence and shock	$\delta = 0.852$, $\sigma = 0.400$	Calibrated by transition-matrix RMSE.
Lead-arrival layer	Feature set	R^Q, R^P, flow index, brands, prior dramas, top-brokerage indicator, matching exposure	Variables governing the structured extensive margin.
Lead-arrival layer	Ridge penalty	10^{-4}	Weak regularization for the arrival-rate logit.
Recursive career	Endorsement-cost grid choice	$\theta_E = 1.0$, $\kappa_E = 25.0$, $\kappa_{HE} = 0.0$	Best grid/SMM choice in the endorsement-capacity extension.
Matching/Bellman closure	Supporting margin	Observed post-cap supporting rate	Held fixed because supporting roles are outside the lead-role matching block.

Notes. Values are the constants used in the current production scripts. Estimated parameters such as ω, φ, a_κ, λ, and $\Delta\kappa_{0p}$ are reported separately in the estimation tables and are not repeated here as tuning parameters.

Maintained assumptions. The main assumptions are: the BLP instruments are excluded from unobserved drama quality after controls; the salary cap changes the bargaining regime rather than the consumer utility of star actors directly; pre-cap competition can be summarized by a scalar Bertrand-intensity parameter; platform-control capability accumulates through prior control experience; the actor law of motion is Markov in (R^Q, R^P) and current role; the matching A-matrix identifies conditional sorting among matched lead pairs; the structured arrival logit summarizes market access, visibility, availability, and search frictions through observed proxies; supporting roles are not explicitly modeled in the matching market; and the actor-career block is partial equilibrium rather than a full market-clearing multi-slot cast equilibrium. These assumptions define what the current model can and cannot claim.

Implementation order. The scripts implement the block structure in the same order as the model. First, the BLP demand block estimates drama demand and converts actor value-added into V_a/R. Second, the bargaining block calibrates ω from the disclosed wage-share moment. Third, the make-or-buy block estimates $(\varphi, a_\kappa, \lambda, \Delta\kappa_{0p})$ from the pre-cap residual moment plus the iQIYI, Tencent, and Youku control-share paths, using 320 random starting points. Fourth, the director-entry block estimates entry costs from the post-cap probit and validates them against the locked-out director DiD. Fifth, the actor-allocation block estimates structured lead arrivals, conditional actor–director sorting, Bellman acceptance, and reputation transitions. Moments mechanically used in estimation are reported as targeted; moments produced by later simulations but not used to estimate that block are validation moments; and facts outside the state space, such as multi-slot cast composition, are reported as tension moments.

5 Estimation Results and Robustness

This section reports the numerical estimates, goodness of fit, and core robustness checks from the model, moment map, and implementation procedure in Section 4. Process details are kept in Section 4.7; here the focus is on what the estimates imply and which modeling features are needed to fit the data.

5.1 Demand

Table 11 reports the BLP demand estimates from Equation (6). Star value-added enters strongly ($\widehat{\beta}_{\text{star}} = 1.42$, $t = 19.8$), as does director value-added ($\widehat{\beta}_{\text{dir}} = 0.89$, $t = 9.0$). The price coefficient is negative and economically sensible ($\widehat{\alpha} = -0.04$, $t = -1.91$). Heterogeneous taste loads on directors ($\widehat{\sigma}_{\text{dir}} = 0.35$) rather than stars, suggesting auteur loyalty rather than a fan-versus-nonfan partition. Genre dummies show narrative dramas valued below romance and period dramas above romance, conditional on cast.

Table 11: BLP Demand Estimates

Parameter	Estimate	SE	t
Constant	−13.85	0.49	−28.5
Price (Hausman IV)	−0.040	0.021	−1.91
log(1 + star VA)	1.422	0.072	19.8
log(1 + director VA)	0.895	0.100	9.0
log(1 + episode count)	1.317	0.078	16.9
Narrative genre	−0.528	0.077	−6.9
Period genre	0.369	0.099	3.7
Post-2019	−1.287	0.117	−11.0
Platform dummies	iQIYI, Tencent, Youku, Mango		
σ_{star}	0.000	0.051	0.0
σ_{dir}	0.353	0.300	1.18

Notes: N = 3,808 dramas in 9 release-year markets. Standard errors in parentheses. Platform dummies and a 2018 year dummy are included but omitted from the table. Random-coefficient standard errors are reported for the unconstrained parameters.

5.2 Platform Nash bargaining

With $V_a/R_j = \beta_{\text{star}}/(1 + \text{raw star VA}) = 0.523$, observed iQIYI base revenue wage share $W/R_j = 0.234$, and calibrated outside-option share $d_S/R_j = 0.04$, the Nash weight solves

$$\widehat{\omega} = \frac{V_a/R - W/R}{V_a/R - d_S/R} = \frac{0.523 - 0.234}{0.523 - 0.040} = 0.599. \tag{21}$$

Platforms therefore capture roughly 60 percent of the bargainable surplus post-cap. The legacy disclosed-deal calibration using $W/R = 0.085$ implies $\omega = 0.907$; broader iQIYI scenario wage shares imply ω as low as 0.369. We use the base revenue median as the central target and report the alternatives as sensitivity.

5.3 Endogenous project-control regime

We estimate Equation (20) by simulated least squares on 24 platform-year control-share observations for iQIYI, Tencent, and Youku from 2017 to 2024, jointly with the pre-cap residual moment that sets the bilateral platform share close to zero before the cap. Table 12 reports the estimates. The Bertrand intensity is $\widehat{\varphi} = 0.998$: in the three-platform fit, pre-cap platform competition pushes the actor's effective outside option close to star MRP. The estimated capability-decay rate is $\widehat{\lambda} = 0.0067$ per year, so platform-control capability is highly persistent once accumulated. The model fits the three control-share paths with RMSE 6.87 percentage points.

Table 12: Make-or-Buy Regime Estimates

Parameter	Estimate	Std. Error	Description
φ	0.998	—	Bertrand outside-option intensity
a_κ	0.908	—	Capability cost-slope in $\Delta\kappa_p(\Psi)$
λ	0.0067	—	Capability decay rate (per year)
$\Delta\kappa_{0,\text{iQIYI}}$	4.580	—	iQIYI baseline make-buy cost wedge
$\Delta\kappa_{0,\text{Tencent}}$	4.523	—	Tencent baseline make-buy cost wedge
$\Delta\kappa_{0,\text{Youku}}$	4.559	—	Youku baseline make-buy cost wedge

Notes: Estimates from simulated least squares on 24 platform-year control-share observations for iQIYI, Tencent, and Youku (2017–2024), plus a pre-cap residual moment $\gamma_M \simeq 0$, using 320 random starting points. The current make-share equation estimates shared φ, a_κ, and λ plus platform-specific $\Delta\kappa_{0p}$. Cost-pass-through variants are reserved for robustness and are not part of the central specification.

5.4 Actor career dynamics

The actor block estimates high persistence in reputation: $\widehat{\delta} = 0.852$ and $\widehat{\sigma} = 0.40$. Table 13 reports the resulting lifetime values from the opportunity-constrained Bellman. The top-tier value $V(H, H) = 152$ million RMB is consistent with industry intuition for an active career of roughly 10 years at average take-home pay of 13–15 million RMB per year. The model fits the post-cap role distribution with RMSE 0.55 percentage points and the transition matrix with RMSE 2.28 percentage points.

Table 13: Actor Lifetime Values $V(s)$ (RMB million, present value)

$R^Q \setminus R^P$	Low	Medium	High
Low	42.5	72.5	87.5
Medium	62.2	98.7	119.5
High	60.6	108.4	152.3

Notes: Values from the actor Bellman with stochastic law of motion ($\delta = 0.852$, $\sigma = 0.40$), structured lead-arrival rates, and matching-generated lead-high/lead-low composition. The recursive closure adds endorsement/capacity choices and remains a partial-equilibrium diagnostic rather than a full cast-market equilibrium.

5.5 Director entry margin

We estimate the binary-choice entry probit of Equation (8) by MLE on the post-cap window (2019–2022). The estimate is $\widehat{\beta}_{\text{Treated}} = -1.012$ (SE 0.071, $z = -14.3$), implying an average marginal effect of −22.6 percentage points: the locked-out cohort enters at 12.6% annually versus 44.6% for streaming incumbents. In κ_d units, locked-out directors have streaming-specific entry costs roughly one standard deviation above the incumbent baseline. This is an incentive-compatibility flip: pre-cap, the Bertrand-driven residual could not cover the narrative non-cast cost, so narrative-type directors stayed out; the cap switches to Nash and expands the residual above the entry cost at the margin.

The two director-entry moments are summarized with the reduced-form evidence in Table 5: the DiD treatment effect on streaming entry is +47.3 percentage points, and the structural supply simulation attributes 2,153 additional director-credits to the cap-induced expansion of the producer residual.

5.6 Matching and model fit

The conditional matching matrix reproduces the 324-cell director–actor cross-tab at RMSE 0.17 percentage points post-cap. Adding the structured extensive-margin arrival layer improves the closed role-choice fit: closed CCP RMSE falls to 5.82 percentage points overall and 5.17 percentage points in cells with at least 50 observations. For the previously problematic high-quality/low-popularity cell, the structured layer implies a 3.39 percent lead-arrival rate and a 1.92 percent accepted-lead rate, close to the 1.63 percent observed rate.

Table 14: Structured Lead-Offer Arrival Rates by Actor State

Actor state	Arrival rate	Accepted lead, model	Accepted lead, data	Gap
H_H	3.02	2.15	1.63	0.52
H_L	3.39	1.92	1.63	0.30
H_M	3.53	2.81	2.23	0.58
L_H	6.17	4.95	4.63	0.32
L_L	4.88	2.87	2.41	0.46
L_M	6.74	5.23	5.32	-0.09
M_H	6.08	5.06	5.13	-0.07
M_L	6.13	3.67	4.39	-0.73
M_M	7.26	5.79	6.24	-0.44

Notes: Entries are percentages from the structured extensive-margin layer. The arrival rate is λ_{it} averaged within each (R^Q, R^P) state. Accepted-lead model rates multiply arrival by the matching-generated high/low composition and Bellman acceptance probabilities. The state-level accepted-lead RMSE is 0.44 percentage points.

5.7 Robustness: ablation checks

The main robustness exercise removes one feature at a time and re-estimates the affected block. The benchmark structured arrival model fits accepted lead rates across actor states with RMSE 0.44 percentage points, the closed role-choice CCP with RMSE 5.17 percentage points in cells with at least 50 observations, and the Bellman role distribution with RMSE 0.55 percentage points. Table 15 reports the corresponding ablations.

The results clarify the economic role of the extensive-margin layer. If lead arrivals depend only on actor quality, or only on quality and popularity, the accepted-lead and closed-CCP fit deteriorates sharply. This rejects the interpretation that good actors mechanically receive the right set of opportunities. The model needs additional market-access variables—visibility, brand exposure, career history, and search exposure—to describe who is actually offered lead roles. The strongest rejection is the matching-only specification: using the conditional matching softmax as the lead-arrival probability makes the high-quality/low-popularity cell mechanically receive nearly universal lead opportunities and raises the Bellman role RMSE to 24.58 percentage points. Conditional sorting is therefore an intensive-margin object; it cannot also serve as the extensive-margin offer-arrival object.

The ablations are also informative about which ingredients are less central. Removing the matching-exposure index alone barely changes the closed-CCP fit once visibility and career-history variables are present. We therefore interpret matching exposure as a useful

Figure 9: Model Fit

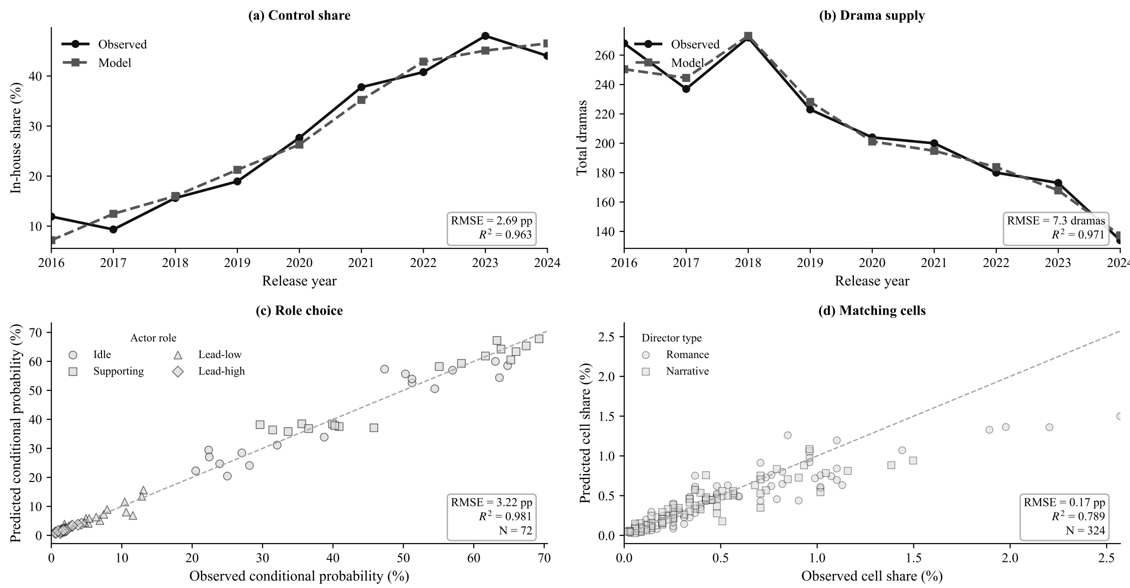

Notes. Panel (a) compares observed and predicted platform-control share; this moment is targeted in the project-control estimation. Panel (b) compares observed and predicted total drama supply; this is an untargeted validation moment. Panel (c) plots observed versus predicted conditional actor role-choice probabilities, with marker shape denoting the actor role. Each point is one state-period-role cell. Panel (d) plots observed versus predicted director–actor matching cell shares, with marker shape denoting director type. Each point is one director-cell by actor-cell pair. Dashed diagonals in panels (c)–(d) are 45-degree lines.

structured proxy for project-side opportunity, not as the sole driver of the arrival layer. Appendix Table 21 reports additional checks on actor acceptance, the stochastic reputation law of motion, popularity as a separate state variable, platform-control dynamics, and SMM starting values. In the updated three-platform make-or-buy block, a static platform-specific wedge fits the raw control-share panel slightly better than the dynamic capability specification; we therefore interpret the dynamic layer as mechanism and counterfactual discipline, while the starting-value check confirms that the 320-start benchmark is not a local-optimum artifact.

Table 15: Main Robustness Checks: Lead-Arrival and Matching Ablations

Ablation	Accepted lead RMSE	Closed CCP RMSE	Bellman role RMSE	Interpretation
Arrival depends only on quality	1.15 pp	11.41 pp	0.86 pp	Lead offers are not allocated only by acting-quality reputation.
Arrival depends on quality and popularity	1.16 pp	11.40 pp	0.86 pp	Two-dimensional reputation is not enough without access and visibility.
No visibility/access variables	0.40 pp	6.22 pp	0.55 pp	Visibility and agency access discipline who receives lead opportunities.
No matching-exposure index	0.47 pp	5.17 pp	0.56 pp	Project-side exposure has limited incremental fit once visibility is included.
Matching-only lead arrival	68.25 pp	37.36 pp	24.58 pp	Conditional sorting cannot by itself explain lead-arrival rates.

Notes. Each row removes one feature from the benchmark model and re-evaluates the relevant fit moments. The main-paper ablations re-estimate the lead-arrival layer where applicable. The appendix ablations report additional mechanism and numerical checks.

6 Results

This section asks whether the estimated model actually supports the proposed mechanism. If the salary cap were just a blunt transfer, the estimates would show wage compression but little evidence that star value-added, bargaining, control rights, and entry line up quantitatively. The estimates say otherwise: stars have high marginal revenue products, the cap lowers the effective outside option, platform residuals rise, the three-platform make-buy block fits the broad control-share transition with visible residual gaps, and the actor-side moments reveal both validation and remaining tension.

6.1 Structural estimates and validation

Table 16 collects the headline structural estimates. The parameters have a simple interpretation. The pre-cap Bertrand intensity is $\widehat{\varphi} = 0.998$: once the pre-cap residual moment is combined with the project-control decision for iQIYI, Tencent, and Youku, the data imply that rival-platform bidding pushed star outside options close to star MRP. The post-cap Nash bargaining weight is $\widehat{\omega} = 0.599$, so the cap restores positive platform surplus rather than transferring the entire surplus to platforms. Capability is highly persistent, with a make-buy wedge decay rate of $\widehat{\lambda} = 0.0067$ per year in the three-platform block.

Table 16: Headline Structural Estimates

Object	Estimate	Source
BLP β_{star}	1.422	Demand (Section 5.1)
BLP α	−0.040	Demand (Section 5.1)
Nash ω	0.599	Post-cap wage shares (Section 5.2)
Bertrand φ	0.998	Pre-cap residual and make-share paths (Section 5.3)
Capability slope a_κ	0.908	Three-platform make-share paths (Section 5.3)
Capability decay λ	0.0067/yr	Three-platform make-share paths (Section 5.3)
Reputation persistence δ	0.852	Actor transitions (Section 5.4)
Transition shock σ	0.40	Actor transitions (Section 5.4)
Locked-out entry gap AME	−22.6 pp	Probit MLE (Section 5.5)

6.2 Bargaining-regime validation

The high star salary is endogenous to scarcity and bargaining, not an exogenous cost shock. From the BLP demand estimates, the marginal revenue product of one lead star has median $V_a/R_j = 0.523$ and a long right tail (Figure 21). Under the pre-cap Bertrand-auction regime, the effective outside option is $d_S^{\text{eff}}/R_j = \widehat{\varphi}V_a/R_j + (1-\widehat{\varphi})d_S/R_j = 0.523$ with $\widehat{\varphi} = 0.998$, so the actor captures $1-\gamma_M = 0.999$ of the star MRP and the per-lead wage is $0.523R_j$ (Figure 8).

The cap severs the rival platform as a credible bidder, collapsing d_S^{eff}/R_j to the non-streaming outside option 0.040. With $\widehat{\omega} = 0.599$, the actor share falls to 0.447, the per-lead wage falls to $0.234R_j$, and the platform residual $\gamma_M V_a$ rises from roughly zero to $0.29R_j$. Equation (20) therefore tilts production toward platform-controlled projects, and Equation (8) admits directors who were previously locked out: the DiD entry effect is +47.3 percentage points and the full simulation attributes 2,153 extra director-credits to the cap (Table 5). The reduced-form tests in Appendix E confirm that stars were not a binding production-function constraint; their high pay reflected bargaining rents, which the cap redirects to the producer side and, through persistent platform-control capability ($1 - \widehat{\lambda} = 0.993$), to variety.

6.3 Structural validation: cast composition

The cast-composition moments in Figure 5 are deliberately held out of the current structural estimation. The model targets the three-platform make-buy paths, validates external-buy volume, and fits actor transition and bilateral matching moments, but it does not yet contain a multi-slot cast choice that would let a platform substitute among top-, mid-, low-tier, and newcomer actors within a drama. Therefore the top-heavy recomposition documented in the empirical section is a validation and tension moment rather than a targeted fit.

Economically, the tension is informative. The bargaining block predicts that the cap lowers the price of top-tier talent relative to its marginal revenue product and restores a positive platform residual. Once top-tier talent is cheaper, platforms can rationally increase the number of top-tier cast slots in surviving tracked dramas, even as total production contracts and mid-tier/newcomer slots disappear. Reproducing the moment structurally would require extending the matching block from one actor-state match per director-project to a portfolio problem with multiple cast slots, tier-specific wages, and an outside-sector option for lower-ranked actors. We therefore use the actor-distribution moment to discipline the next model extension rather than to estimate the current baseline.

6.4 Structural validation: actor outside option

The matched Endata–Douban endorsement panel provides a second held-out actor moment. Table 7 follows the same actors from 2017–2018 to 2019–2022 and shows that high pre-cap rating-reputation actors lose high-quality lead roles (−2.75 pp), gain low-quality lead roles (+2.07 pp), and increase matched endorsement contracts from 0.06 to 0.33. Unmatched actor-years are retained as zero matched contracts, so the level should be read as a conservative matched-sample measure rather than a population endorsement rate. Within this high-reputation bin, the correlation between endorsement growth and high-quality lead growth is negative, $\rho = -0.182$, whereas the corresponding correlations for low- and mid-reputation actors are positive.

This pattern supports the model's treatment of endorsements as part of the actor outside option, but it is not yet a targeted Bellman moment. In the current baseline, endorsement income is a passive flow $E_{at}(R^Q, R^P)$ and the actor chooses only one streaming role tier. Matching the endorsement substitution moment structurally would require an actor capacity constraint that allocates time across high-quality leads, low-quality leads, endorsements, variety, and idling, with a time-varying commercial outside option. We therefore use endorsements as an external validation of the outside-option channel and as a diagnostic for the next actor-block extension.

7 Counterfactuals

This section asks what would have happened without the cap and which part of the policy matters for welfare. If the cap only shifted rents, consumer surplus and production variety would be unchanged once revenue is held fixed. The model's accounting block does show a large transfer from actors to platforms, but the production block shows why the transfer is economically consequential: the restored platform residual finances project control, persistent control capability, director entry, and quality-adjusted variety.

7.1 Surplus decomposition

The full simulation combines a dynamic project-control module, a competitive external-content module, and a budget module. The welfare proxy is quality-adjusted drama count. The updated three-platform project-control block matches the iQIYI, Tencent, and Youku control-share paths with RMSE 6.87 percentage points. The downstream supply simulation matches the total drama count with RMSE 7.3 dramas (about 3% of the mean), despite the latter being an untargeted moment.

Table 17 reports the headline surplus accounting for 2019–2022. Under the cap, roughly RMB 181 billion is shifted out of the actor wage stack and into platform profit in the fixed-revenue accounting block. This table is therefore a transfer decomposition, not a complete welfare estimate. The welfare-relevant margin is the production response that follows from the restored platform residual: project control, persistent platform capability, director entry, and quality-adjusted variety. Appendix Figure 25 visualizes the no-cap trajectory for control share, director entry payoff, producer profit gap, and actor value.

Table 17: Counterfactual Surplus Decomposition, 2019–2022 (billion RMB)

Scenario	Consumer Surplus	Drama Revenue	Actor Wages	Platform Profit
Factual (cap in place)	120.2	313.1	146.5	−111.8
No cap	120.2	313.1	327.4	−292.7
Δ (cap minus no-cap)	0.0	0.0	−180.9	+180.9

Notes: The no-cap counterfactual holds consumer surplus and drama revenue fixed at their factual levels and resets the platform bargaining share to the pre-cap Bertrand level. The RMB 180.9 billion transfer is the actor-wage accounting effect.

Figure 10 ties these effects to the model primitives. Panel (a) decomposes the 2017–2024 rise in platform-control share using the dynamic project-control equation (20): the rise reflects the external-supply shock, persistent platform-control capability, and their interaction. Panel (b) reports the stage-1 surplus accounting: the cap compresses actor gross wages by roughly RMB 181 billion, directors gain RMB 4.5 billion from expanded entry, and consumer surplus is unchanged in the transfer-only counterfactual.

Figure 10: Model-based counterfactual decomposition. Panel (a): sources of the rise in platform-control share, 2017–2024. Panel (b): stage-1 surplus transfers under the cap versus the no-cap counterfactual, 2019–2022.

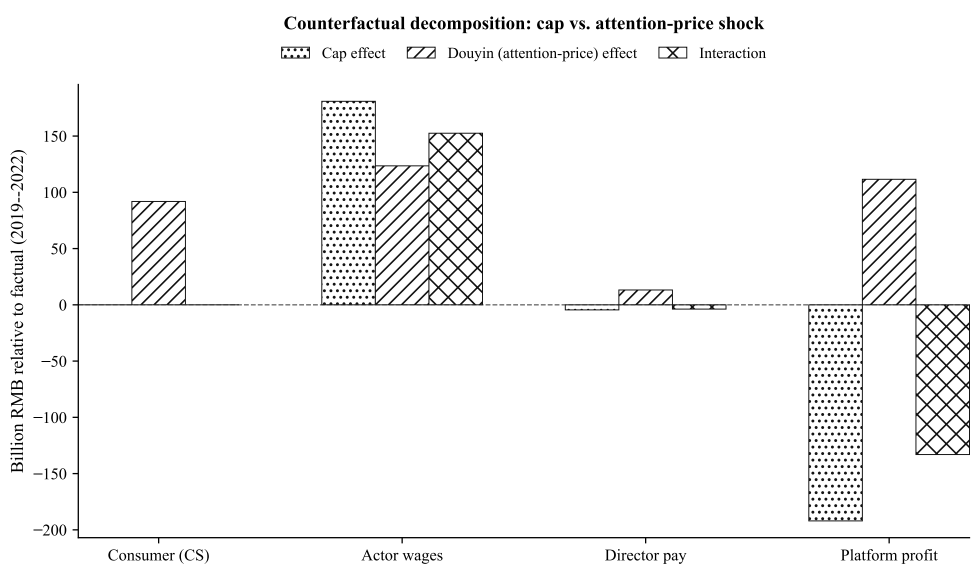

Figure 11 summarizes the transmission mechanism. The cap is a public commitment that removes the rival platform as a credible bidder, lowering the actor's effective outside option d_S^{eff} and restoring a positive platform Nash residual. The resulting wage compression reduces actor gross wages by roughly RMB 181 billion in the current accounting block. The enlarged platform residual finances platform-controlled production, while accumulated platform-control capability is highly persistent ($1-\widehat{\lambda} = 0.993$ per year). That same residual admits roughly 2,153 additional director-credits (+RMB 45 billion), while the relative decline in external-studio surplus ($V_k - c_k$) reduces the number of bought dramas. The net effect is more quality-adjusted and director-entry variety, even though the total tracked drama count falls because the external-supply contraction dominates on the quantity margin.

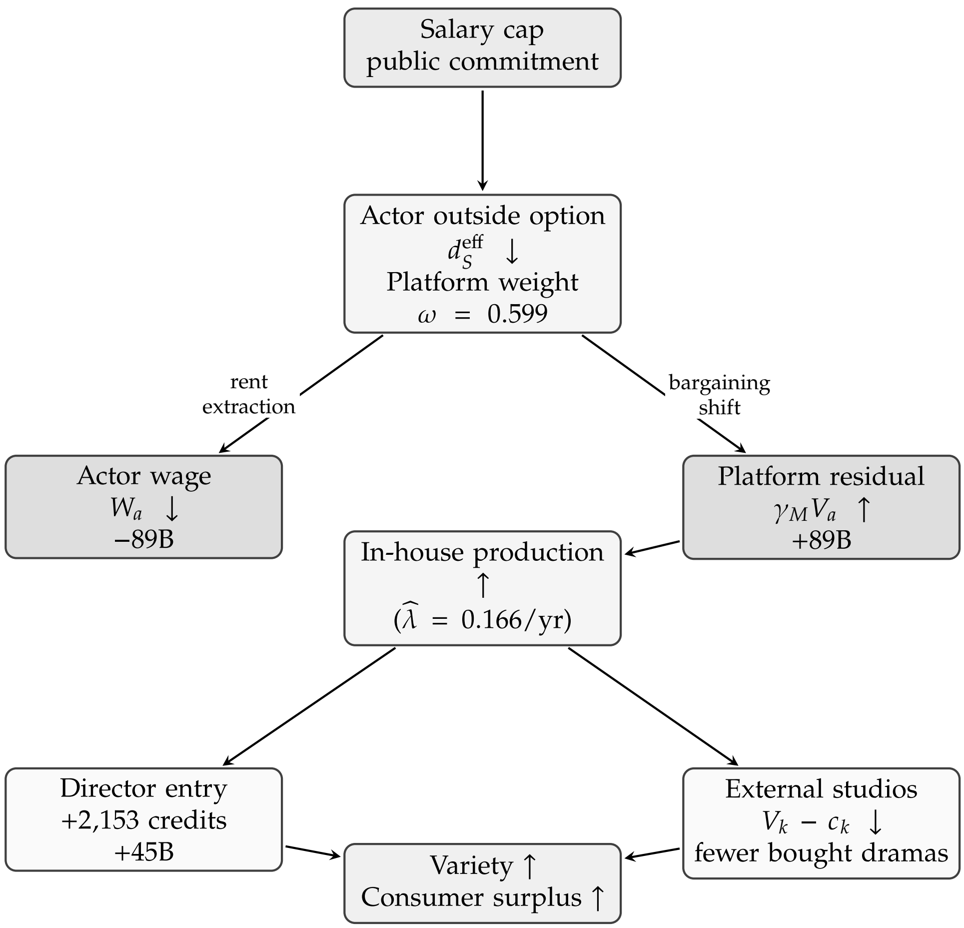

Figure 11: Transmission of the salary cap: from the bargaining-regime switch to the reallocation of surplus and the production response.

7.2 Timing and no-cap counterfactuals

Table 18 collects the timing and extension counterfactuals. An early cap in 2018 is worse than the actual 2019 cap because the supply-shock loss from early studio exit outweighs the small platform-control gain. A late cap in 2020 is better because delaying studio exit keeps more external dramas in the market. The timing exercises should be read as scenario accounting around the downstream supply module: external-studio exit dominates total quantity, while the make-or-buy block mainly reallocates the composition of controlled versus externally sourced projects.

7.3 Attention price and the Douyin shock

The rise of short-form video is an exogenous shock to the price of long-form attention. We model this as an attention-price index

$$P_t = \frac{\text{long-form share of total video time}_t}{\text{long-form share of total video time}_{2017}}, \tag{22}$$

so that $P_{2017} = 1$ and Douyin's entry drives P_t down to 0.44 by 2022 (Table 18). Revenue, consumer surplus, and star wages in year t are scaled by P_t relative to the no-Douyin benchmark, while production costs are held fixed. The long-form share is constructed from calibrated industry-report center points as long-form user-minutes divided by total long- plus short-form user-minutes, rather than taken as a directly reported statistic.

Table 18 reports the resulting 2×2 decomposition for 2019–2022. The salary cap compresses actor wages by roughly RMB 181 billion; removing it would lower platform profit by roughly RMB 192 billion in this accounting, while leaving consumer surplus and total drama revenue unchanged. The attention-price shock is larger on revenue: holding the cap in place but restoring P_t to its 2017 level raises drama revenue by RMB 264 billion, consumer surplus by RMB 92 billion, and actor wages by RMB 124 billion, while platform profits rise by RMB 112 billion.

The cap is primarily a transfer between actors and platforms; the attention-price shock is a level effect that alters every margin. Both forces jointly compressed actors—the no-cap-no-Douyin counterfactual would have delivered actor wages of roughly RMB 603 billion versus the factual RMB 147 billion—but only the cap restored platform surplus and variety.

7.4 Endogenous external-studio exit

The baseline simulation treats the number of external (licensed) dramas J_{buy} as exogenous. In practice, external studios exit when licensing becomes unprofitable. We endogenize

Table 18: Counterfactual Extensions

Panel A: Timing scenarios, deviations from actual-cap baseline

Scenario	ΔJ_{buy} avg/yr	ΔJ_{make} avg/yr	ΔJ_{total} avg/yr	ΔW cumulative
Early cap (2018)	−22	+0.9	−16.2	−203
Late cap (2020)	+14	+0.5	+12.1	+160
No cap	0	−0.9	−1.2	−14
Structural γ_M reset	0	−7.2	−8.9	−104

Panel B: Salary-cap and attention-price decomposition, 2019–2022 totals (billion RMB)

Outcome	Factual	No cap	No Douyin	No cap & no Douyin	Cap effect	Douyin effect
Consumer surplus	120.2	120.2	212.1	212.1	+0.0	+91.9
Drama revenue	313.1	313.1	577.0	577.0	+0.0	+264.0
Actor wages	146.5	327.4	270.1	603.5	+180.9	+123.5
Director pay	15.7	11.1	28.9	20.5	-4.5	+13.2
Platform profit	-111.8	-303.8	-0.3	-325.3	-192.0	+111.6

Panel C: Attention-price path and endogenous attention outside value

Year	Long-form share	P_t	Year	Predicted long share	No-Douyin CF
2017	86.2%	1.000	2017	0.851	0.851
2018	64.6%	0.749	2018	0.688	0.874
2019	54.2%	0.629	2019	0.544	0.885
2020	50.3%	0.583	2020	0.461	0.910
2021	45.0%	0.522	2021	0.450	0.933
2022	37.9%	0.440	2022	0.393	0.938

Panel D: Endogenous external-studio exit

	Pre-cap (2016–2018)	Post-cap (2019–2024)
Quality-cost index $\log(Q/c)$	12.235	11.709
Observed J_{buy}	228	121
Predicted J_{buy}	227	121
No-cap counterfactual J_{buy}	228	228
Cost-quality frontier β	0.152	
Cap shift γ	-0.526	
Supply elasticity η	1.21	

Notes. J_{buy} denotes externally controlled dramas and J_{make} platform-controlled dramas. P_t is the long-form share of total video time normalized to 2017. This share is constructed from calibrated industry-report center points as long-form user-minutes divided by total long- plus short-form user-minutes. The endogenous attention model is $\log(s_L/s_S) = \alpha + \beta_S \log(q_S) + \beta_M \cdot make\_share$, with $\widehat{\beta}_S = -3.93$, $\widehat{\beta}_M = 3.10$, and $R^2 = 0.978$. Panel D calibrates $\log J_{\text{buy}} = \alpha + \eta \log(Q/c)$ to match pre/post averages.

this margin by combining the p5c cost-quality frontier with a free-entry supply curve.

The p5c regression estimates

$$\log(\text{views}_k) = \beta \log(c_k) + \gamma \cdot \text{post} + \text{genre FE} + \varepsilon_k, \tag{23}$$

with $\widehat{\beta} = 0.152$ and $\widehat{\gamma} = -0.526$. Because $\widehat{\beta} < 1$, the salary cap raises external-drama quality by less than it lowers cost, so the quality-cost wedge $\log(Q/c)$ falls from 12.235 pre-cap to 11.709 post-cap. A free-entry curve

$$\log J_{\text{buy},t} = \alpha + \eta \log(Q/c)_t, \tag{24}$$

calibrated to the observed pre- and post-cap averages, yields $\widehat{\eta} = 1.21$. The implied endogenous path tracks the observed average well: 228 external dramas pre-cap and 121 post-cap (Table 18). The no-cap counterfactual holds the wedge at its pre-cap level, so J_{buy} would have stayed near 228.

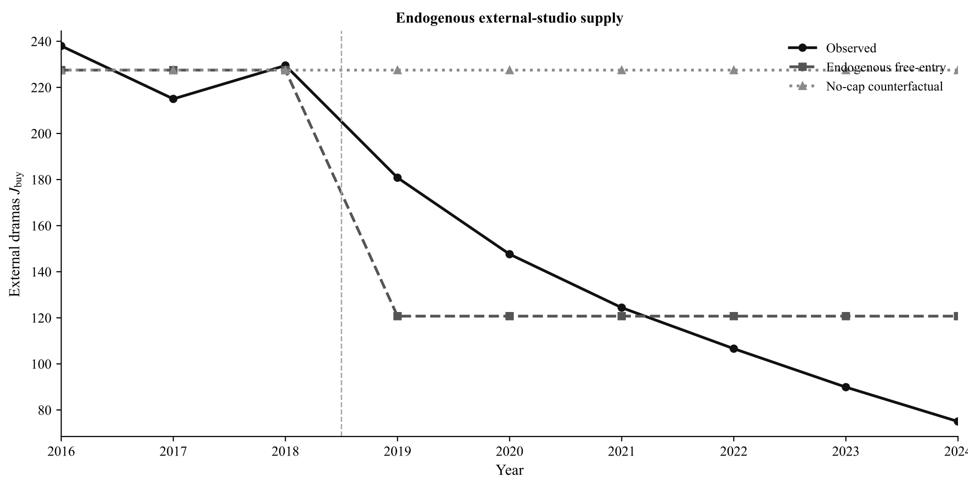

Figure 12: Endogenous external-studio exit. The solid line is the observed number of bought dramas; the dashed line is the free-entry prediction; the dotted line is the no-cap counterfactual.

7.5 Endogenizing the attention outside option

The preceding counterfactuals treat the attention-price index P_t as an exogenous shock driven by Douyin. We now model P_t as the outcome of a viewer choice between long-form and short-form video. Let s_{Lt} be the long-form share of total video time and q_{St} short-form minutes per user day. We estimate

$$\log\left(\frac{s_{Lt}}{1-s_{Lt}}\right) = \alpha + \beta_S \log(q_{St}) + \beta_M \cdot \text{make\_share}_t + \varepsilon_t, \tag{25}$$

using 2017–2022 macro attention data. The estimates are $\widehat{\beta}_S = -3.93$ and $\widehat{\beta}_M = 3.10$ ($R^2 = 0.978$). Douyin growth is the dominant force pulling attention away from long-form video, while the rise in platform-control share partially offsets it.

Table 18 reports the no-Douyin counterfactual: if q_{St} had remained at its 2017 level, the long-form share would have stayed near 85–94% rather than falling to 38%. The endogenous attention-price index $P_t = s_{Lt}/s_{L,2017}$ therefore attributes almost the entire decline in long-form attention to short-form competition, not to a fall in long-form quality.

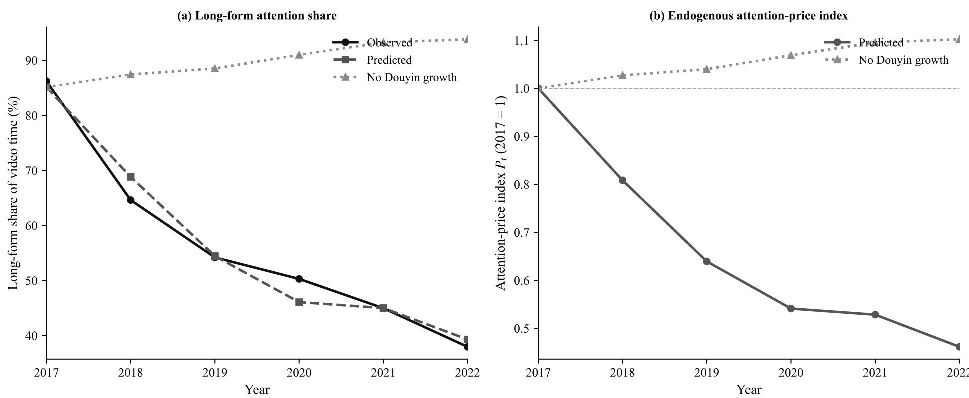

Figure 13: Endogenous attention outside value. Panel (a): observed and predicted long-form attention share, and the no-Douyin counterfactual. Panel (b): predicted attention-price index with and without Douyin growth.

8 Conclusion

The concluding question is what the China episode teaches beyond a single labor regulation. If the lesson were only that wage caps lower wages, the contribution would be narrow. The evidence instead shows that platforms do not merely match upstream producers to downstream consumers. They transmit their monetization metric backward into the production process. When attention is the metric that platforms can monetize, competition for scarce attention-producing inputs can turn downstream attention rents into upstream input rents and narrow the set of feasible products.

The 2018 Chinese actor salary cap functioned not as a marginal labor-cost intervention but as a public commitment device in an attention-funded platform market. By rendering a binding ceiling on above-line compensation enforceable, the policy weakened the Bertrand-auction loop between traffic stars, advertising revenue, subscriber acquisition, and the production-cost floor that had narrowed the pre-cap production route. Structurally, the pre-cap residual moment and three-platform make-or-buy block imply high pre-cap Bertrand intensity ($\widehat{\varphi} = 0.998$); the cap reduced the actor outside option, restored positive platform bargaining power ($\widehat{\omega} = 0.599$), and admitted directors who had been incentive-incompatible with the prior regime. Entry-margin dramas account for a large share of post-cap consumer-side utility under the estimated BLP demand; the accounting counterfactual implies roughly RMB 181 billion in actor-wage compression over 2019–2022.

Generalization. The mechanism is not specific to China. Wherever an upstream input has tournament-shaped rewards and downstream platforms monetize a single attention metric, competition can capitalize that reward into bargaining rents and create production monoculture. The distortion is therefore not only on price, entry, ranking, recommendation, or the consumer side. It can appear as an upstream input tournament that determines what kinds of products are feasible. The closest analogues are other long-form streaming and recorded-media markets, where revenue models reward concentrated attention and platform competition raises the price of scarce talent (Aguiar and Waldfogel, 2018; Waldfogel, 2018). The Chinese cap is one realization of a broader regulatory category—public input-price commitment in attention-funded platform markets—that, to our knowledge, has not been previously identified in the empirical industrial-organization literature.

Limitations. The model abstracts from the single-tour-per-actor capacity constraint, from year-to-year variation in outside options driven by brand-market cycles, and from a fully structural market-clearing matching model. Embedding per-impression ad revenue separately from subscription revenue, and allowing time-varying star marginal utility as short-form video grows, are natural next steps.

References

ACEMOGLU, D. AND S. JOHNSON (2023): *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, New York: PublicAffairs.

ADLER, M. (1985): "Stardom and Talent," *American Economic Review*, 75, 208–212.

AGUIAR, L. AND J. WALDFOGEL (2018): "Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music," *Journal of Political Economy*, 126, 492–524.

ANDERSON, S. P. AND S. COATE (2005): "Market provision of broadcasting: A welfare analysis," *The review of Economic studies*, 72, 947–972.

ARCIDIACONO, P. AND R. A. MILLER (2011): "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 79, 1823–1867.

ARMSTRONG, M. (2006): "Competition in Two-Sided Markets," *RAND Journal of Economics*, 37, 668–691.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile prices in market equilibrium," *Econometrica*, 63, 841–890.

BERRY, S. T. AND J. WALDFOGEL (1999): "Free Entry and Social Inefficiency in Radio Broadcasting," *Quarterly Journal of Economics*, 114, 397–420.

COOPER, R., D. V. DEJONG, R. FORSYTHE, AND T. W. ROSS (1990): "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80, 218–233.

DESTA, Y. (2018): "Where Has Chinese Mega-Star Fan Bingbing Disappeared To?" `https://www.vanityfair.com/hollywood/2018/09/fan-bingbing-tax-evasion-scandal-disappearance`, accessed July 1, 2026.

FRANK, R. H. AND P. J. COOK (1995): *The Winner-Take-All Society: How More and More Americans Compete for Ever Fewer and Bigger Prizes, Encouraging Economic Waste, Income Inequality, and an Impoverished Cultural Life*, New York: The Free Press.

GREENWALD, B. C. AND J. E. STIGLITZ (1986): "Externalities in Economies with Imperfect Information and Incomplete Markets," *Quarterly Journal of Economics*, 101, 229–264.

HO, K. AND R. S. LEE (2017): "Insurer Competition in Health Care Markets," *Econometrica*, 85, 379–417.

Hotz, V. J. and R. A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60, 497–529.

iQIYI, Inc. (2018–2024): "Annual Reports on Form 20-F," `https://www.sec.gov/edgar/browse/?CIK=1722608&owner=exclude`, revenue, cost-of-revenue, and content-cost line items compiled from Item 5 of annual Form 20-F filings; accessed July 1, 2026.

Mankiw, N. G. and M. D. Whinston (1986): "Free Entry and Social Inefficiency," *RAND Journal of Economics*, 17, 48–58.

Meixler, E. (2018): "Missing Chinese Actress Fan Bingbing Fined $70 Million in Mysterious Tax Evasion Case," `https://time.com/5413538/fan-bingbing-fined-china-tax-evasion/`, accessed July 1, 2026.

National Radio and Television Administration (2021): "National Radio and Television Administration Statements on Illegal and Non-Compliant Performers and Institutions," `https://www.nrta.gov.cn/`, statements reported by state media on August 27, 2021; accessed July 1, 2026.

Nevo, A. (2001): "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 69, 307–342.

Rochet, J.-C. and J. Tirole (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1, 990–1029.

Rosen, S. (1981): "The economics of superstars," *The American Economic Review*, 71, 845–858.

Spence, A. M. (1976): "Product Selection, Fixed Costs, and Monopolistic Competition," *Review of Economic Studies*, 43, 217–235.

State Taxation Administration of China (2021): "Tax Administrative Penalty in the Zheng Shuang Case," `https://www.chinatax.gov.cn/`, reported by state media on August 27, 2021; accessed July 1, 2026.

Waldfogel, J. (2003): "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," *RAND Journal of Economics*, 34, 557–568.

——— (2018): *Digital Renaissance: What Data and Economics Tell Us about the Future of Popular Culture*, Princeton University Press.

Yuan, J. (2019): "The Untold Story Behind the Mysterious Disappearance of Fan Bingbing," `https://www.vanityfair.com/hollywood/2019/03/the-untold-story-disappearance-of-fan-bingbing-worlds-biggest-movie-star`, accessed July 1, 2026.

Appendices

A Equilibrium Formalization

This appendix asks when a salary cap changes production rather than only splitting surplus differently. If pre-cap bargaining already left producers a positive residual, the cap should not change the feasible set of director types. The model shows the opposite case: when Bertrand bidding pins $W^* = V_a$ and $\gamma_M = 0$, producer surplus is $S = R - C_{\text{prod}} - V_a$. For idol production, $C_{\text{prod}}(\text{idol})$ is low and R is high because the majority taste is large; production is feasible. For narrative production, $C_{\text{prod}}(\text{narrative}) = \bar{c}$ is high and the audience is smaller; $S < 0$ at the margin, so narrative directors are less likely to enter. The equilibrium is therefore tilted toward traffic-star formats rather than technology-neutrally covering all formats.

Post-cap, the Nash wage is $W^* = (1-\omega)V_a + \omega d_S < V_a$ because $d_S < V_a$ and $\omega > 0$. The platform share is $\gamma_M = \omega(1 - d_S/V_a) > 0$. Producer surplus becomes $S = R - C_{\text{prod}} - W^* > R - C_{\text{prod}} - V_a$; the residual non-cast budget expands. For narrative production, $S > 0$ at the margin once $\bar{B} - W^* > \bar{c}$, admitting narrative directors. The equilibrium becomes multi-content. Capability accumulation lowers each platform's effective project-control cost through $\Delta\kappa_p(\Psi_{pt}) = \Delta\kappa_{0p}\exp(-a_\kappa \Psi_{pt})$, while λ governs depreciation of the capability stock Ψ_{pt}, shifting more production toward platform control.

B Demand Details

This appendix asks how the demand estimates turn drama characteristics into star marginal revenue product. If demand were estimated only from raw popularity, the bargaining block would inherit mechanical star effects. The BLP procedure instead uses market shares, price instruments, and leave-one-out competitor instruments. For each characteristic $x_{jt}^{(k)}$ that enters mean utility, we compute the leave-one-out competitor sum $Z_{c,jt}^{(1)} = \sum_{j' \neq j, j' \in t} v_{c,j't}$ and the squared deviation from the market mean $Z_{c,jt}^{(2)} = (v_{c,jt} - \bar{v}_{c,-j,t})^2$. The Hausman price instrument is the mean subscription price of competing platforms in the same year. The contraction mapping solves $\delta^{(n+1)} = \delta^{(n)} + \log s_{jt}^{\text{obs}} - \log s_{jt}(\delta^{(n)}, \sigma)$ until convergence.

C Actor Bellman Details

This appendix asks how much actor career choice can be captured without solving a full market-clearing cast equilibrium. If actors freely chose roles, observed lead-role rates would directly reveal preferences; if roles are matching outcomes, opportunity probabilities also matter. The outside-option calibration sets $d_S(L, L) = 0.5$ million RMB and $d_S(H, H) = 30$ million RMB pre-cap, with intermediate cells interpolated by a multiplicative form $d_S(s) = \text{anchor}_L \cdot (R^P(s)/R^P_L)^\eta \cdot q(R^Q(s))$ with $\eta = 1.5$ and quality premium $q(L), q(M), q(H) = 1.0, 1.5, 2.0$. Post-cap, the calibration is refined to per-state direct values to match the role distribution; the supporting wage anchor is raised to 5 million RMB to reflect cameo and art-film compensation. The recursive extension iterates between Bellman values, actor acceptance probabilities, and opportunity composition until convergence, holding total post-cap role opportunities fixed.

D Matching Details

This appendix asks whether the model can reproduce who matches with whom after the cap. If matching were irrelevant, actor career dynamics could be estimated independently of the director side. The data instead require an allocation layer, estimated with the conditional logit likelihood $\mathcal{L}(A) = \sum_{(i,j_{\text{obs}},t)} \log \Pr(j_{\text{obs}} \mid i, t)$, where $\Pr(j \mid i, t) = \exp(\Phi(x_i, y_j))/\sum_{j' \in C(t)} \exp(\Phi(x_i, y_{j'}))$. The actor features are R^Q, R^P, log cumulative views, age, and log traffic index; director features are mean Douban quality and genre shares. Separate A matrices are estimated pre-cap ($N = 503$ matches) and post-cap ($N = 3{,}913$ matches).

E Pre-Cap Star Dependence Tests

This appendix asks whether the pre-cap market was technologically dependent on traffic stars. If stars were a hard production-function constraint, reducing their wages would mainly distort quality. Five reduced-form tests reject that strong dependence. (i) Year-by-year log-view regressions yield pre-cap mean $\beta_\star = 1.50$ versus post-cap mean $\beta_\star = 1.63$. (ii) A hit-rate logit shows the probability that a mid-star drama reaches the top 20% of views rises from 10.5% to 17.8%. (iii) Only 12% of top-20% pre-cap dramas featured a top-5% star. (iv) Across 64 specification curves, only 10.9% show pre-cap $\beta_\star$ exceeding post-cap. (v) Pooling all years and interacting inputs with a post-cap dummy yields positive and significant interactions for both star VA and director VA, although the writer VA interaction is null. The evidence supports the Bertrand-auction interpretation captured by $\widehat{\varphi} = 0.998$ in the pre-cap residual and three-platform make-or-buy block.

F Identification Map

This appendix table records the mapping from empirical moments to structural objects. It is the table version of the prose moment map in Section 4.7: the main text explains the logic, while Table 19 provides the audit trail for which moment identifies, calibrates, validates, or leaves tension for each component.

Table 19: Empirical Moments, Identified Objects, and Estimation Status

Block	Empirical moment	Parameter or object disciplined	Status
Demand	Drama market shares by release year and platform subscription-price variation	BLP mean utility, price coefficient α, product-characteristic coefficients β, random-coefficient terms Σ, and drama unobservables ξ_j	Upstream input; not counted as a structural SMM moment
Attention	Long-form video attention share falls from 86.2% in 2017 to 37.9% in 2022	Attention outside-value path used in counterfactual revenue environment	Targeted/calibrated path
Bargaining	Disclosed post-cap iQIYI wage share, W/R = 23.4%	Nash bargaining weight ω given V_a/R and calibrated d_S/R	Calibrated
Make-or-buy	Pre-cap platform residual plus iQIYI, Tencent, and Youku platform-control shares by platform-year, 2017–2024	Bertrand intensity φ, capability cost-slope a_κ, capability-depreciation rate λ, and platform-specific baseline control-cost wedges $\Delta\kappa_{0p}$	Targeted by SMM/NFXP
Actor Bellman	Post-cap actor role distribution by reputation state	Actor flow-payoff calibration, acceptance probabilities, and role-choice values	Targeted
Actor law of motion	9-cell year-over-year actor reputation transition matrix	Reputation persistence δ and transition shock σ	Targeted
Lead arrival	Accepted lead rates by actor state, combined with Bellman acceptance probabilities	Structured offer-arrival coefficients η governing the extensive margin λ_{at}	Targeted
Matching	Director-type by actor-state match-cell shares	Conditional sorting matrix A and lead-high versus lead-low composition conditional on arrival	Targeted
Director entry	Post-cap streaming-entry probit and locked-out director DiD	Director entry-cost distribution and residual-surplus entry response	Targeted and validated
Industry supply	Yearly total drama count and external-buy volume	Downstream supply simulation implied by project-control and residual-surplus path	Validation
Cast composition	Top-star, mid-tier, and newcomer cast-share shifts	Multi-slot cast substitution margin not in baseline state space	External/tension
Actor endorsements	Endorsement contracts rise for high-reputation actors while high-quality leads fall	Outside-option and capacity-substitution interpretation in recursive actor extension	Validation/tension

Notes. The BLP demand row is reported because it supplies V_a/R and demand-side primitives to the structural pipeline, but its product-share moments are not counted as structural SMM moments. Targeted structural moments enter estimation directly. Calibrated moments pin down a parameter by construction. Validation moments are not used to estimate the listed object and therefore discipline the model out of target. External/tension moments document empirical patterns that require an extension beyond the current baseline, such as a full multi-slot cast-choice problem.

Table 20: Robustness Checks for the Identification Evidence

Outcome	Specification	Coef. (SE)	N	Note
iQIYI gamma	post dummy	12.078* (6.376)	7	Annual T=7.
iQIYI gamma	2018 to 2022	31.609	7	-8.6% to 23.0%.
iQIYI gamma	pre-trend extrapolation	43.511	7	Pre-trend predicts -20.5%; actual is 23.0%.
iQIYI gamma	gamma_hat_pct post + t	-6.250 (14.325)	7	Post dummy with linear trend; underpowered annual series.
iQIYI gamma	gamma_hat_pct P_t	-27.254** (13.653)	7	Attention-threat diagnostic.
iQIYI gamma	gamma_hat_pct dau_m	0.031** (0.016)	6	Douyin-threat diagnostic.
Residual placebo	iQIYI vs Netflix, platform/year FE	6.901 (6.731)	14	Tests whether iQIYI residual recovery exceeds a global streamer benchmark.
iQIYI audited make share	post dummy	0.239*** (0.049)	9	Audited 20-F content-cost flows.
iQIYI audited make share	2018 to 2024	0.284	9	15.6% to 44.0%.
iQIYI audited make share	post + trend	0.012 (0.055)	9	Dynamic rise is absorbed by trend.
Three-platform control share	platform FE, clustered by platform	0.206*** (0.031)	24	Only 3 clusters; use as descriptive panel check.
Three-platform control share	platform FE + trend	-0.026 (0.053)	24	Post jump is not separate from dynamic ramp-up.
Three-platform control share	exclude Tencent	0.191*** (0.041)	16	Leave-one-platform-out.
Three-platform control share	exclude Youku	0.234*** (0.032)	16	Leave-one-platform-out.
Three-platform control share	exclude iQIYI	0.191*** (0.042)	16	Leave-one-platform-out.
Director entry	main cohort, year FE, clustered director	0.681*** (0.030)	5225	Treatment is no 2018 streaming presence.
Director entry	Established directors only (>=5 pre-2017 credits).	0.712*** (0.033)	3560	Established directors only (>=5 pre-2017 credits).
Director entry	Exclude immediate 2019 transition from post definition.	0.398*** (0.038)	5225	Exclude immediate 2019 transition from post definition.
Director entry	Exclude 2022.	0.636*** (0.032)	4180	Exclude 2022.
Douban score	year + platform + genre FE	0.117** (0.057)	2394	Project-level validation outcome.
log 30-day views	year + platform + genre FE	0.884*** (0.082)	3808	Attention outcome.
Douban score	non-imputed sourcing only	0.160 (0.098)	819	Drops calibrated imputation.
log 30-day views	non-imputed sourcing only	0.161* (0.088)	931	Drops calibrated imputation.
Annual mean Douban score	pre/post annual mean	0.423	9	5.79 to 6.21; validation, not main identification.
Attention index P_t	P_t on Douyin DAU	-0.001*** (0.000)	6	Threat diagnostic: attention falls sharply as Douyin grows.

Notes. The annual residual and audited make-share series are short, so trend-controlled specifications are reported as diagnostics rather than preferred estimates. Director-entry standard errors are clustered by director. Three-platform control-share specifications have only three platform clusters and should be read descriptively.

G Robustness

This appendix reports five additional ablation checks that are useful for interpretation but less central than the lead-arrival checks in Section 5.7. The first two remove actor-side decision or transition structure. Forcing actors to accept every lead opportunity worsens the role-distribution RMSE, and using the deterministic reputation law of motion worsens the actor fit relative to the stochastic transition. The third check removes popularity as a state variable by returning to the quality-only arrival specification; the accepted-lead RMSE rises from 0.44 to 1.15 percentage points. The fourth check removes dynamic platform-control capability from the make-or-buy block; in the three-platform panel, this improves raw control-share fit, so the dynamic layer should be read as mechanism and counterfactual discipline rather than as a feature forced by the control-share RMSE alone. The final check varies the number of SMM starting points for the three-platform make-or-buy block.

Table 21: Appendix Robustness Checks: Career, Project-Control, and Numerical Ablations

Ablation	Fit statistic	Benchmark	Ablated	Interpretation
No acceptance margin	Actor role RMSE	0.55 pp	1.06 pp	Actors do not mechanically accept every lead opportunity.
Deterministic reputation transition	Actor role RMSE	0.55 pp	0.89 pp	Career evolution needs stochastic reputation updating.
No popularity state	State accepted-lead RMSE	0.44 pp	1.15 pp	Popularity is a separate state, not noise around quality.
No dynamic platform capability	Make-or-buy RMSE	6.87 pp	5.24 pp	Control-share fit alone does not force the dynamic capability layer.
Fewer SMM start points	Make-or-buy RMSE range	6.87 pp	6.87–6.87 pp	The 320-start benchmark is not driven by a local optimum.

Notes. Each row removes one feature from the benchmark model and re-evaluates the relevant fit moments. The main-paper ablations re-estimate the lead-arrival layer where applicable. The appendix ablations report additional mechanism and numerical checks.

H Additional Evidence

This appendix asks whether the empirical and structural claims survive when the supporting moments are shown explicitly. If the main text were cherry-picking a few attractive figures, the appendix would reveal weak or disconnected diagnostics. The organizing principle here is evidentiary rather than mechanical: the first block documents the reduced-form facts behind the market transition, the second block reports model-fit and robustness checks, and the remaining blocks collect compact tables for moments that are useful for replication and interpretation.

H.1 Reduced-Form Evidence

The figures in this subsection show why the salary-cap episode should be read as a change in the production environment rather than as a narrow wage-accounting rule. Figure 14 places long-form streaming inside the broader attention-market squeeze from short-form video. Figures 15–18 document the reduced-form patterns that motivate the model: rewards are concentrated around traffic inputs; producer participation broadens; and genre composition shifts rather than expanding uniformly. Figures 19 and 20 provide external market context for the advertising and international comparison margins.

Figure 14: Long-form squeeze: Douyin DAU versus iQIYI subscribers, 2017–2022.

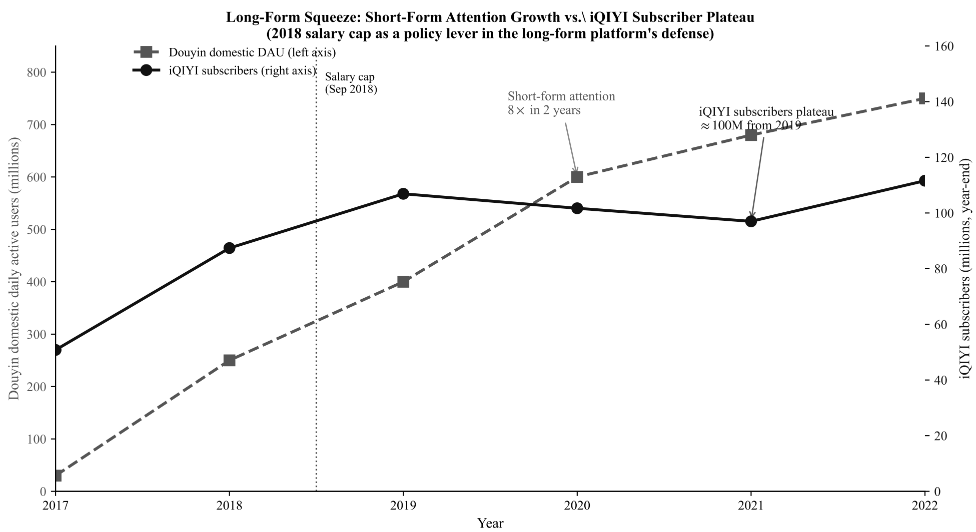

Figure 15: Convexity tests. Panel (a): production-side returns to star input, pre-cap dramas. Panel (b): mean log reward by within-year rank decile.

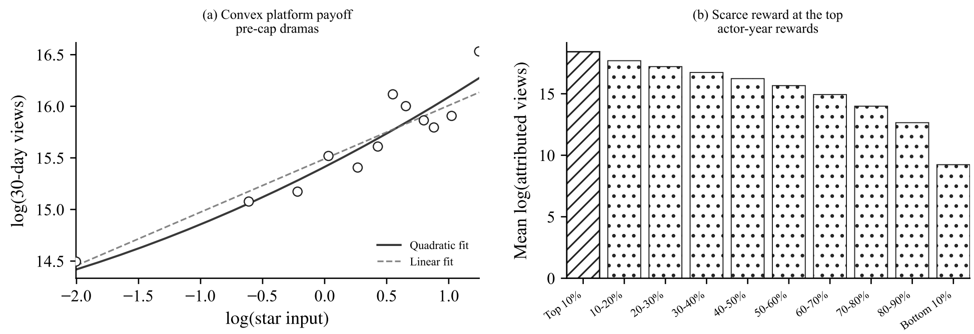

Figure 16: Variety diagnostics: producer participation, category coverage, and quality dispersion.

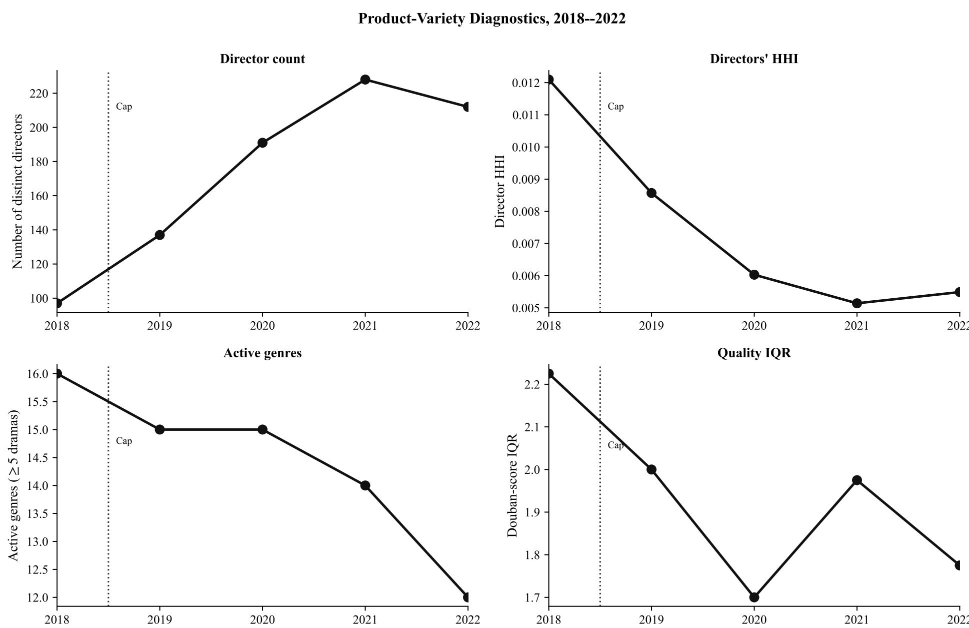

Figure 17: Genre shift pre- versus post-cap.

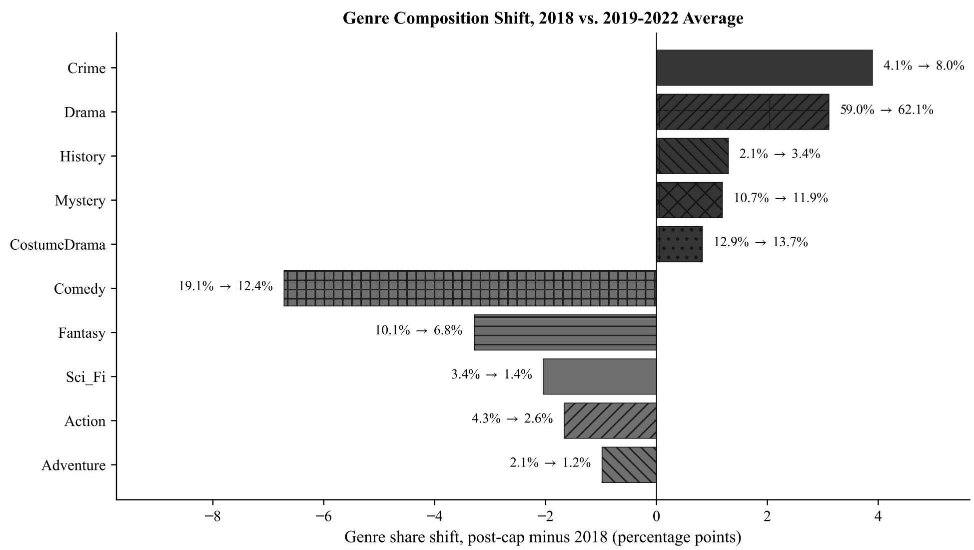

Figure 18: Evolution of genre shares, 2017–2025.

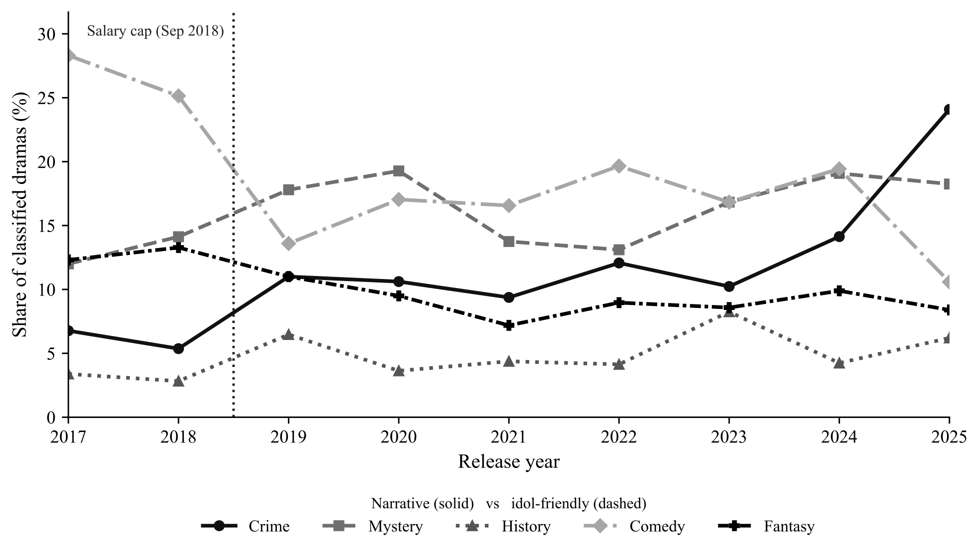

Table 22: Genre Composition Shift, Pre vs. Post

Genre	Pre (2018)	Post (avg)	Δ (pp)
Top gainers			
Crime	4.1%	8.0%	+3.90
Drama	59.0%	62.1%	+3.11
History	2.1%	3.4%	+1.30
Mystery	10.7%	11.9%	+1.19
CostumeDrama	12.9%	13.7%	+0.83
Top losers			
Comedy	19.1%	12.4%	-6.72
Fantasy	10.1%	6.8%	-3.28
Sci_Fi	3.4%	1.4%	-2.04
Action	4.3%	2.6%	-1.67
Adventure	2.1%	1.2%	-0.98

Notes: Within-year drama share by Douban genre, 2018 vs. 2019–2022 average. Shares sum to more than 100% because dramas carry multiple genre tags. Ordered by percentage-point change; top five gainers and losers shown.

Reading: The post-cap supply shifts from spectacle / star-vehicle genres (Comedy, Fantasy, Sci-Fi, Action) toward narrative-driven genres (Crime, History, Mystery, Drama, Costume Drama). The pattern is consistent with the two-front squeeze defense: long-form streaming differentiates against short-form by investing in narrative depth, the production technology of which short-form formats cannot replicate. Comedy alone falls 6.7 percentage points (a −35% relative decline); Crime nearly doubles its share.

Table 23: Genre Shift Decomposition: Entry vs. Incumbent Transition

Genre	Incumbent pre 2018	Incumbent post avg	Entry post avg	Total shift (pp)	Entry contrib. (pp)	Transition contrib. (pp)
Top gainers						
Crime	1.7%	4.2%	8.7%	+5.81	+5.12	+0.69
Mystery	7.5%	8.5%	13.5%	+4.67	+4.40	+0.27
History	2.5%	3.3%	4.7%	+1.81	+1.59	+0.22
Action	1.7%	1.9%	3.8%	+1.63	+1.57	+0.06
Drama	66.7%	65.6%	67.1%	+0.00	+0.30	-0.30
Top losers						
Comedy	16.7%	7.5%	8.7%	-8.30	-5.85	-2.45
Romance	45.8%	50.0%	34.8%	-6.92	-8.04	+1.12
Fantasy	10.8%	9.0%	6.2%	-3.86	-3.36	-0.50
Sci_Fi	5.0%	0.5%	1.4%	-3.86	-2.64	-1.22
Family	7.5%	2.4%	5.5%	-2.81	-1.43	-1.38

Notes: Decomposition of the post-cap genre composition shift into (i) the entry-margin contribution — entry directors having a different genre type than incumbents — and (ii) the incumbent-transition contribution — incumbent directors shifting their own genre mix post-cap. "Incumbent" = directors with $\geq$ 1 streaming drama in 2018; "Entry" = directors with 0 streaming dramas in 2018 but $\geq$ 1 during 2019–2022 and $\geq$ 2 pre-2017 directing credits (the treated cohort of Section 3).
Counterfactual "no transition" post share: incumbents hold the 2018 genre mix and only entry directors add their observed genre composition. The entry-contribution column is the share change under this counterfactual; the transition-contribution column is the residual.
Reading: The genre shift is overwhelmingly driven by the entry-margin composition channel, not by incumbent transition. For the top gaining genres, entry directors account for 88–94% of the total share gain (Crime +5.12 of +5.81pp, Mystery +4.40 of +4.67pp, History +1.59 of +1.81pp). For the losing genres, the picture is similar: 70% of the Comedy decline is the entry margin (entry directors carry a far lower Comedy share than incumbents had pre-cap), and the entire Romance decline (−8.04pp from entry alone) is offset by incumbents actually shifting *into* Romance (+1.12pp transition). The substantive interpretation is that the post-cap genre mix reflects *who* produces the new content (a different cohort of directors specializing in different genres) rather than incumbents reinventing themselves.

H.2 Model Validation

This subsection reports diagnostics for the structural objects used in the counterfactuals. Figure 21 shows that star marginal revenue products are high across the post-cap distribution, supporting the bargaining interpretation rather than a story driven only by a few outliers. Figures 22–24 check downstream payoffs, actor dynamics, and robustness to taste-shock dispersion. Figure 25 summarizes the no-cap paths used to translate the estimated mechanism into counterfactual project-control, entry, and surplus accounting.

Figure 19: Advertising-market bifurcation: long-tail share versus blockbuster tail.

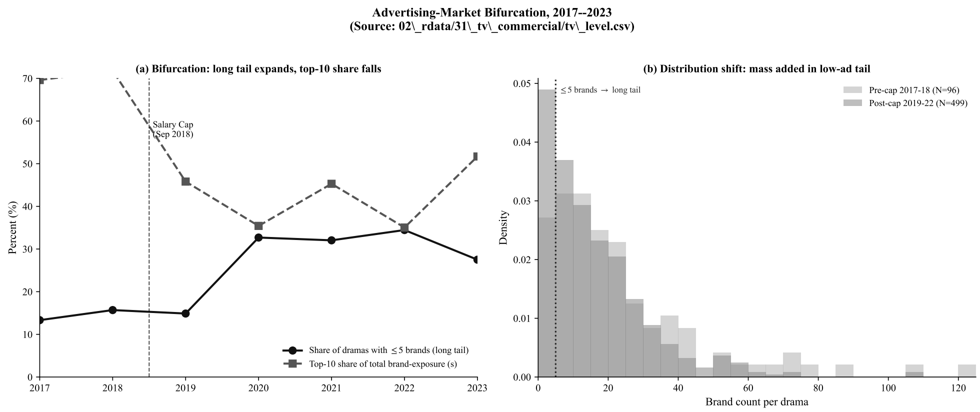

Figure 20: US/Netflix overlay: comparative cap-free benchmark.

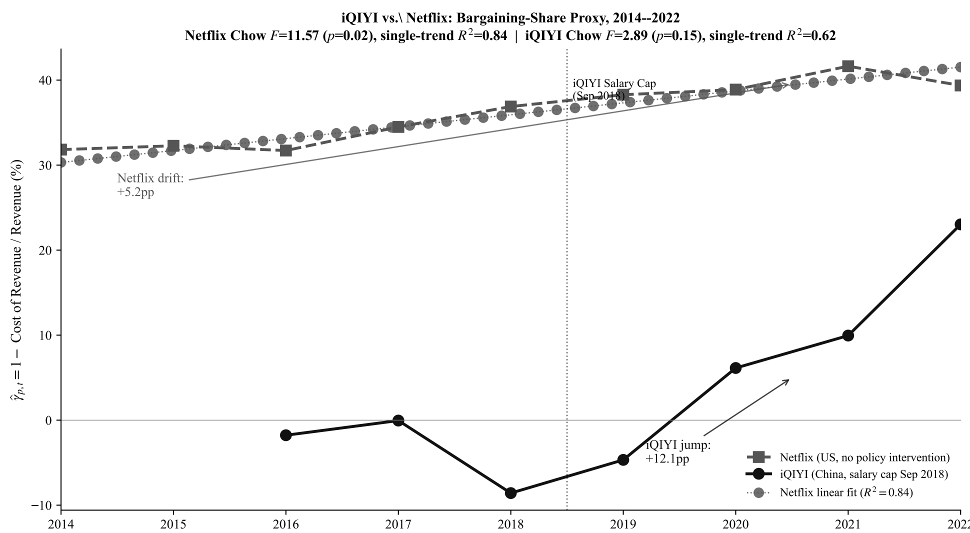

Figure 21: Distribution of the BLP-implied star marginal revenue product V_a/R_j across post-cap dramas. Panel (a): density with median, 90th percentile, post-cap Nash wage share, and outside option. Panel (b): cumulative distribution function.

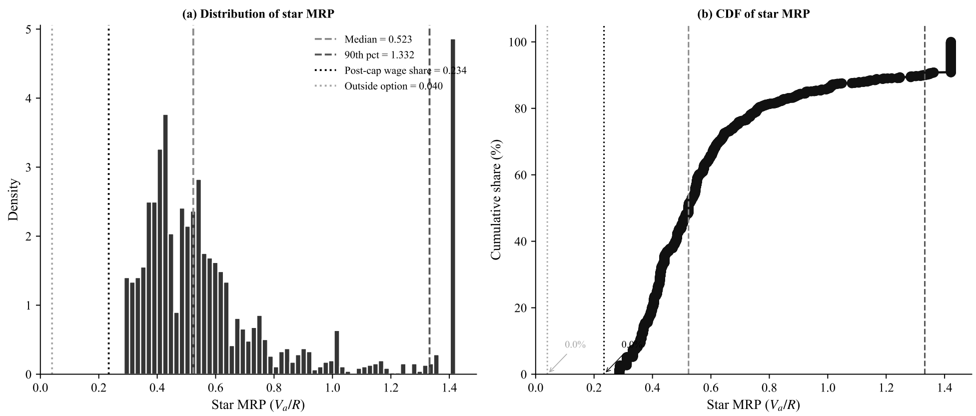

Figure 22: Downstream audit: regime-aware wages and payoffs.

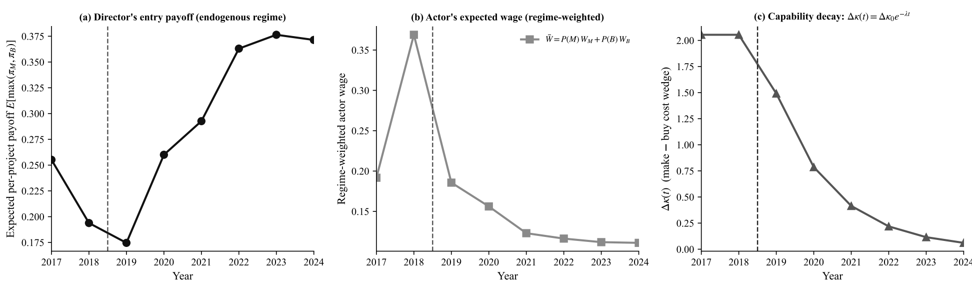

Figure 23: Actor dynamic parameters: value functions and CCPs.

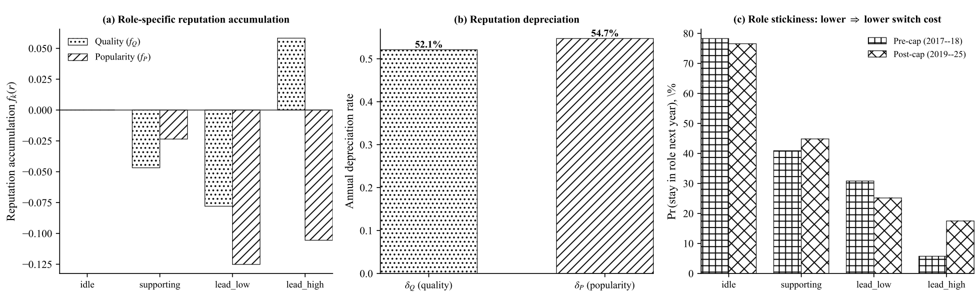

Figure 24: Regime-choice taste-shock robustness grid.

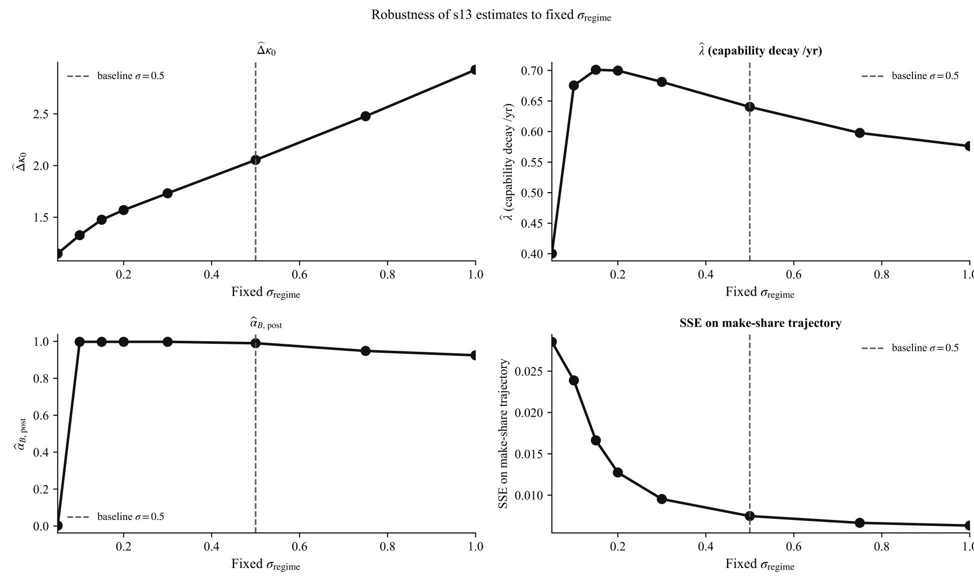

Figure 25: Structural no-cap counterfactual trajectories.

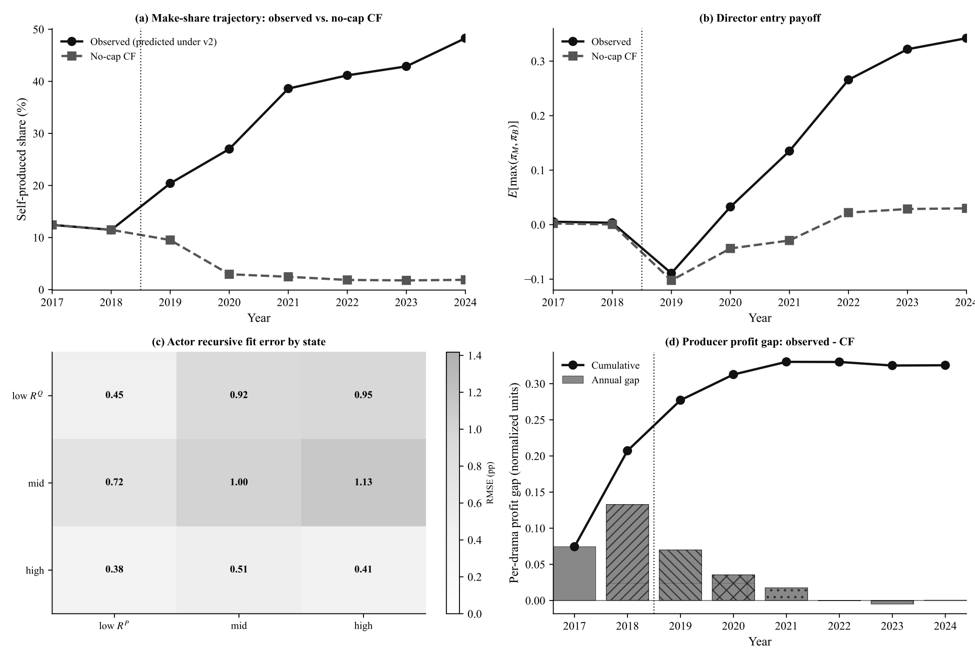

H.3 Empirical Diagnostics

Table 24 collects the empirical moments that are too detailed for the main text: sample summary statistics, convexity estimates, project-control quality gaps, the full star-MRP distribution, and compact model-fit diagnostics. The table is meant to make the mapping from descriptive facts to targeted structural moments auditable in one place.

H.4 Actor Outside Options and Capacity

Table 25 documents how the actor-side career data enter the model. The endorsement panel is sparse, so the table reports coverage explicitly before using it as an outside-option proxy. The static and recursive capacity panels then show whether the career-choice extension can reproduce the observed lead-role and endorsement margins without treating actors as freely choosing roles independently of market opportunities.

H.5 Platform, Advertising, and Entry Margins

Table 28 gathers the platform-side moments that connect the empirical facts to the structural mechanism. The project-level residual panel motivates why restored producer residuals matter. The advertising panel supports the interpretation that a single traffic-reward system concentrates attention. The pipeline and director-origin panels show the extensive-margin response: platform control expands the set of leads and draws in directors who were active outside the pre-cap streaming core. Table 27 separates released-title supply from pipeline and sponsorship entry, clarifying why a tournament can produce many attempts while realized attention remains concentrated.

H.6 Counterfactual and Welfare Accounting

Table 29 separates accounting transfers from welfare-relevant production responses. Panel A reports how much platform operating income changes under alternative no-cap residual assumptions. Panel B uses the demand estimates to quantify the consumer-surplus contribution of entry-margin dramas, which is the welfare channel behind the headline wage-transfer accounting.

Table 24: Appendix Empirical and Model-Fit Diagnostics

Panel A: Core sample summary

Measure	N	Mean	Std. dev.	Min	Max
Douban score	2,394	6.168	1.432	2.100	9.600
Log(views, 30-day)	3,808	14.161	2.574	8.517	19.296
Episode count	3,822	30.141	13.469	1.000	160.000
Actor attention index	44,867	1.965	1.328	0.033	12.965
Director VA, views	28,921	0.813	1.536	-6.943	3.962
Actor VA, views	36,597	0.400	1.394	-6.812	4.583

Panel B: Convexity tests

Estimate	Coefficient	Robust SE
Reward side: log(rank)	6.243	0.229
Reward side: $[\log(\text{rank})]^2$	-0.618	0.016
Production side: log(star input)	0.617	0.077
Production side: $[\log(\text{star input})]^2$	0.060	0.033

Panel C: Project-control quality and attention

Measure	Externally controlled	Platform controlled
Matched dramas	2,601	1,221
Rated dramas	1,430	964
Mean Douban score	6.051	6.342
Share score ≥ 7	0.165	0.247
Mean log 30-day views	13.928	14.659
Controlled score gap	0.117 (0.057), $p = 0.042$	
Controlled view gap	0.884 (0.082), $p < 0.001$	

Panel D: Star marginal revenue product distribution

Statistic	V_a/R	Statistic	V_a/R
Mean	0.628	Median	0.523
25th percentile	0.416	75th percentile	0.688
90th percentile	1.332	95th percentile	1.422
Share > post-cap wage share	100.0%	Share > outside option	100.0%

Panel E: Model-fit diagnostics

Moment	Model	Observed / interpretation
BLP observables explain log-share variance	$\widehat{\beta}' X_j$	$R^2 = 0.358$ on $N = 1{,}580$
Random-coefficient placement	$\sigma_{\text{dir}} > \sigma_{\text{star}}$	$\widehat{\sigma}_{dir} = 0.153$, $\widehat{\sigma}_{star} = 0.0002$
Director HHI	Falls post-cap	0.012 → 0.006
Top-tercile actor career path	lead_high → lead_low	lead_high 5.3% → 3.3%; lead_low 4.0% → 8.7%

Notes. Panels A–D report the main data diagnostics: sample moments, reward/production convexity, project-control quality gaps, and star-MRP distribution. Panel E collects the high-level model-fit checks formerly reported separately. Project-control gaps include platform, release-year, and genre fixed effects. The MRP calculation uses $V_a/R = \beta_{\text{star}}/(1+a)$ over post-cap dramas with $\beta_{\text{star}} = 1.422$.

Table 25: Actor Outside-Option and Capacity Diagnostics

Panel A: Endorsement-panel coverage

Measure	Value	Share
Actor-year rows in state panel	93,063	100.0%
Unique actors in state panel	16,700	100.0%
Actors with at least one matched endorsement	872	5.2%
Actor-years with at least one endorsement	2,648	2.8%
Mean endorsement contracts per actor-year	0.09	—

Panel B: Selected endorsement coverage by pre-period actor state

Rating	Popularity	Actors	Matched history	Endorsed actor-years	Mean contracts
H	H	2,159	4.9%	1.6%	0.041
L	H	2,894	9.6%	4.5%	0.162
L	L	3,060	5.6%	1.2%	0.025
M	H	4,047	8.9%	4.9%	0.212
M	M	3,771	10.5%	3.7%	0.106

Panel C: Static career-capacity benchmark

Rating bin	ΔP(lead high)	ΔP(lead low)	Δ endorsements	$\rho(\Delta E, \Delta LH)$
Low	-2.31	+0.30	+0.76	+0.034
Mid	-1.26	+0.28	+0.46	+0.094
High	-0.83	-0.44	+0.21	+0.086

Panel D: Dynamic and recursive closure diagnostics

Exercise	Rating	Data LH	Model LH	Data endorse.	Model endorse.
Dynamic capacity	Avg. L/M/H	1.7%	1.7%	2.6%	1.4%
Recursive matching	Avg. L/M/H	1.7%	2.4%	2.6%	1.4%
Dynamic capacity	RMSE		0.67 pp		
Recursive matching	RMSE		0.83 pp		

Notes. Endorsements are matched to the Douban actor panel using the name bridge from `p0b_actor_salary_proxy.py`; unmatched actor-years are retained as zero contracts, so moments are conservative matched-sample moments. State labels cross rating and popularity terciles; Panel B reports the diagonal and high-popularity cells to keep the appendix table compact. The static benchmark uses pre/post changes in the unified actor-year panel. The dynamic capacity grid has $(\theta_E, \kappa_E, \kappa_{HE}) = (1.0, 25.0, 0.0)$. The recursive matching loop converges in 13 iterations and is a partial-equilibrium closure, not a full market-clearing cast model.

Table 26: External Evidence for the Star-Tournament Background

Background premise	External evidence	Source	Role in the paper
High star pay compressed production budgets	The 2018 salary ceilings tied total cast compensation to production cost and limited the lead actor's share of the cast wage bill. Contemporary coverage connects the rule to "sky-high" compensation, split contracts, and tax enforcement.	Meixler (2018); Desta (2018); Yuan (2019)	Establishes that extreme star compensation was an institutional problem, not an artifact of our data.
Hidden compensation contracts were economically meaningful	The Fan Bingbing case involved alleged dual contracts; the Zheng Shuang case involved reported agreed compensation of RMB 160 million for one drama and follow-on tax and broadcast sanctions.	Meixler (2018); State Taxation Administration of China (2021); National Radio and Television Administration (2021)	Supports the interpretation that observed wage shares can understate true pre-cap compensation pressure.
Platform monetization rewarded attention	iQIYI Form 20-F filings report online advertising as a major revenue line. In the compiled revenue panel, online advertising averages 45 percent of iQIYI revenue during 2016–2018, before falling to 22 percent during 2019–2024 as membership becomes larger.	iQIYI, Inc. (2018–2024); `p0h_iqiyi_revenue.csv`	Motivates why traffic stars enter the platform's willingness-to-pay rather than only affecting artistic quality.
The tournament itself must be measured in data	External records do not observe the bidding game. The paper measures it using pre-cap star input, 30-day views, brand sponsorship concentration, actor attention shares, and pipeline-entry diagnostics.	Figure 4; Table 6; Appendix Tables 27 and 28	Provides empirical moments that the structural model must address.

Notes. "Star tournament" is the paper's economic label for a market in which platforms bid for scarce actors who convert concentrated attention into bargaining rents. The external sources document the institutional background; the paper's own data identify the mechanism.

Table 27: Quantity, Concentration, and Lottery-Ticket Entry

Evidence	Pre / baseline	Post / comparison	Reading
Main released-title panel: Tracked dramas	466 in 2018	400 in 2024	The market-wide released-title panel does not show a simple post-cap quantity boom; output becomes more selected.
Actor attention rewards: Top-10 actor share of attributed 30-day views	47.1%	annual range 43.7%–50.6%	Attention remains winner-take-most even when many projects are produced.
Advertising validation: Sponsor-data drama count	51 in 2018	125 per year, 2019–2022	The externally observed sponsorship sample expands below the blockbuster tier.
Advertising validation: Top-10 brand-exposure share	71.8%	40.4%	More sponsored titles coexist with concentrated advertising payoff.
Advertising validation: Long-tail sponsored titles	15.7%	28.5%	The extra mass is mainly low-sponsorship long-tail projects, not a collapse of the top tier.
iQIYI pipeline validation: Production announcements	32, 2015–2018	367, 2019–2025	Platform-side pipeline entry rises sharply; this is the lottery-ticket margin, not necessarily released-title supply.
iQIYI pipeline validation: Unique announced leads	120 lead-year appearances	1296 lead-year appearances	Pipeline expansion draws in many potential leads while attention remains concentrated ex post.

Notes. Released-title counts use the Yunhe–Douban drama panel. Actor attention rewards use attributed 30-day views from the actor-drama VA panel. Advertising validation uses Endata drama-level brand sponsorship records. Pipeline validation uses iQIYI public production announcements. The table separates market-wide released supply from platform pipeline entry: a winner-take-most tournament can generate many attempts to create hits while realized attention and sponsorship remain concentrated at the top.

Table 28: Platform Residual, Advertising, and Entry-Margin Diagnostics

Panel A: Per-drama γ_p^{drama}, iQIYI internal panel

N (disclosed cost & revenue)	15
10th percentile	-18.869
25th percentile	-5.301
Median	**-2.180**
75th percentile	-1.346
90th percentile	-0.065
Maximum	+0.572
Firm-level $\widehat{\gamma}_p$, 2018	−0.086
Firm-level $\widehat{\gamma}_p$, 2022	+0.230

Panel B: Advertising-market bifurcation

Period/year	N	Median brands	Long tail	Top tier	Gini	Top-10 share
2018	51	15.0	15.7%	15.7%	0.530	71.8%
2019–2022 avg.	125	12.9	28.5%	4.1%	0.460	40.4%
2023	120	12.0	27.5%	0.8%	0.440	51.7%

Panel C: iQIYI pipeline expansion

Period	Annc.	Lead slots	Unique leads	Leads/drama	New leads
2015–2018	32	128	111	4.00	111
Post-cap pool, 2019–2025	367	1,459	765	3.98	687
Post/pre ratio	11.5×	11.4×	6.9×	—	6.2×

Panel D: Origin of post-policy director entrants

Career stage at first credit	N	Share
Truly new	32	8.0%
Recent (2015–16 debut)	34	8.5%
Veteran (2000–09)	125	31.2%
Legacy (pre-2000)	150	37.4%
Average pre-2017 Douban score, post-only	2.374	
Average pre-2017 Douban score, both-period	2.494	
Post-2018 high-quality dramas led by post-only directors	81.9%	

Notes. Panel A uses realized iQIYI projects with disclosed cost and revenue; project-level medians are unweighted, while the firm-level benchmark is revenue-weighted. In Panel B, long tail denotes dramas with at most five sponsoring brands, top tier denotes at least 50 brands, and top-10 share is the concentration of brand exposure on the ten largest dramas each year. Panel C uses iQIYI public production announcements; the post-cap pool contains 765 unique leads, of whom 687 had not appeared in the pre-cap iQIYI lead pool. Panel D defines post-only directors as directors credited on 2019–2022 dramas but absent from the 2018 streaming core; the 2017 streaming-drama panel has incomplete director identifiers.

Table 29: Counterfactual and Welfare Accounting

Panel A: Legacy iQIYI accounting sensitivity, cumulative 2019–2022 gap (RMB millions)

Revenue counterfactual \ γ_p^{CF}	0	−0.086	−0.035
Realized revenue	+10,187	+20,357	+14,326
Membership frozen at 2018 level	+10,187	+18,391	+13,526
All revenue lines frozen at 2018 level	+10,187	+18,783	+13,686

Panel B: Consumer-surplus entry-margin contribution

Year	Total dramas	Entry-margin dramas	CS observed	CS no-entry	Entry-margin share
2019	154	93	0.150	0.058	61.5%
2020	203	135	0.165	0.058	64.8%
2021	228	137	0.202	0.102	49.4%
2022	223	139	0.220	0.078	64.7%
Avg. 2019–2022	—	—	—	—	60.1%

Notes. Panel A is a legacy iQIYI-only operating-income sensitivity, not the three-platform structural make-or-buy estimate. It reports $\sum_t(\Pi_t^{actual} - \widetilde{\Pi}_t)$ under alternative revenue and no-cap accounting-margin assumptions, with SG&A and R&D held at realized values; positive entries mean the cap raises cumulative operating income relative to no cap. Panel B computes per-market consumer surplus as the log-sum-exp over inside utility in the estimated BLP demand model; the no-entry benchmark removes dramas directed by the locked-out cohort.