Regulating the Traffic Economy: Salary Caps and Vertical Integration in Streaming*

Yu Hao †
The University of Hong Kong

Yele Ma ‡
The University of Hong Kong

May 31, 2026
Preliminary — please do not circulate

Abstract

Ad-funded digital platforms systematically under-supply variety relative to heterogeneous consumer demand. We show that in talent-intensive platform markets, this preference externality [Anderson and Coate, 2005, Berry and Waldfogel, 1999, Waldfogel, 2003] is amplified by an upstream Bertrand auction for scarce stars that drives talent wages to their full marginal product and forecloses variety-feasible production technologies. China's 2018 actor salary cap provides a quasi-experimental shock that breaks the auction. On a panel of 3,822 streaming dramas, we estimate the pre-cap Bertrand outside-option intensity at $\widehat{\varphi} = 0.998$ and the cap-triggered capability accumulation rate at $\widehat{\lambda} = 0.64/\text{yr}$. The cap-induced variety expansion contributes 60% of post-cap consumer-side log-sum-exp utility; a structural counterfactual yields RMB 10–20 billion in averted iQIYI losses. The welfare gain operates primarily through consumer variety, not through wage redistribution—a Kaldor-Hicks improvement in a tournament-input platform market.

*Preliminary version. **Acknowledgements and disclosure.** We gratefully acknowledge financial support from the Hong Kong University Grants Committee's Collaborative Research Fund (CRF), under the project "Algorithmic Bias, Economic Efficiency, and Social Welfare" (Project Coordinator: Yanhui Wu; grant number on file with the funding agency). We thank Ke Wang at iQIYI for assistance with access to industry data used in this study; iQIYI did not review, fund, or otherwise influence the analysis, framing, or conclusions of this paper, and was not given pre-publication review rights. Yele Ma worked as a consultant for iQIYI during her PhD studies, prior to the inception of this project, and has no ongoing financial or contractual relationship with the firm. Yu Hao has no financial relationship with iQIYI or any other Chinese long-form video streaming platform. The authors have no other relevant financial interests, paid or unpaid positions, or relationships that could be perceived as a conflict of interest. We are grateful for comments and discussions from Yanhui Wu and from colleagues at the HKU Business School.

†yhao@hku.hk Faculty of Business and Economics, The University of Hong Kong.

‡yelema@connect.hku.hk Faculty of Business and Economics, The University of Hong Kong.

Keywords: Preference Externality, Two-Sided Platforms, Tournament Inputs, Salary Cap, Structural Estimation, Variety Welfare, Bertrand-Talent Auction.

JEL Codes: L13, L82, L86, J33, K23, D43.

1 Introduction

Why does ad-funded streaming under-supply variety? 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 the minority taste is systematically underweighted. This paper argues that in talent-intensive platform markets, the preference externality is amplified by an upstream Bertrand auction for scarce stars whose welfare consequences are first-order and structurally identifiable. We show this empirically using a quasi-experimental policy shock—China's 2018 actor salary cap—in the world's largest streaming market by audience size, with annual revenue exceeding RMB 80 billion across the three major platforms by 2018. The Chinese setting is unusually clean: a public, enforceable, dated regulatory intervention from a major emerging-market regulator that breaks the auction without otherwise altering the underlying demand-side environment.

The mechanism in three steps. First, under Bertrand competition for scarce traffic stars, the wage is bid up to the star's full value-added: $W = V_a$. Second, this collapses Nash bargaining between platform and talent ($\gamma_M = \omega(1 - V_a/V_a) = 0$): the platform extracts zero residual surplus from any in-house project, so all production routes to the buy regime where the producer is an external studio. Third, since the external studio's budget is exhausted by the star wage, only formats that justify a top-star investment (idol/traffic dramas serving the majority taste) are feasible; narrative- or quality-led variants are foreclosed by budget arithmetic. The Nash equilibrium is a unique monoculture even when every party recognizes it as inefficient: five reinforcing forces hold the equilibrium in place—an ad-revenue technology that loads concentration on the top decile of star input; procurement-side risk aversion that tilts platform buyers toward bankable cast [Holmström, 1979]; a platform-level subscriber-acquisition race that subordinates single-project ROI to long-run market-share growth [Sutton, 1991]; superstar tournament dynamics that make unilateral exit unilateral surrender [Rosen, 1981, Adler, 1985, Frank and Cook, 1995]; and the budget-arithmetic exclusion that renders narrative production infeasible once the star wage absorbs the residual. The cap operates on the last force directly and unwinds the other four because they were sustained by the same Nash equilibrium with no enforceable exit.

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). The three major platforms (iQIYI, Tencent Video, Youku) 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 build a structural model that takes the consumer-demand side from Berry et al. [1995], Nevo [2001] BLP random-coefficients, the talent- market wage formation from Nash bargaining [Ho and Lee, 2017], and a project-level endogenous make-or-buy regime that closes the producer's choice problem. Five primitives identify the equilibrium: the BLP preference coefficients, the Nash bargaining weight ω, the Bertrand outside-option intensity φ, the make-buy cost wedge $\Delta\kappa$, and the cap-triggered capability-accumulation rate λ. We estimate them sequentially on a panel of 3,822 streaming dramas (2017–2025) using ω identified from $N = 10$ disclosed iQIYI projects, φ and $\Delta\kappa$ from the year-by-year make-share trajectory, and the actor career-path Bellman from a Hotz- Miller two-step CCP procedure.

Headline estimates and their welfare implications are reported in the abstract and developed in detail in Sections 5 and 6.

Contributions and connection to the literature. First, we extend the preference-externality framework of Anderson and Coate [2005] and the free-entry-inefficiency results of Spence [1976], Mankiw and Whinston [1986], Berry and Waldfogel [1999] to a setting where the variety bias is amplified by an upstream tournament-input mechanism absent from the standard two-sided platform framework of Rochet and Tirole [2003], Armstrong [2006]. Second, we connect the superstar-rent literature [Rosen, 1981, Adler, 1985, Frank and Cook, 1995] and the sports salary-cap literature [Dietl et al., 2012, Kahn, 2007] to the welfare implications of content-platform wage regulation, quantified on a 3,822-drama panel [Berry et al., 2019]. Third, we provide the first structural identification of φ and λ as separately identified primitives in a make-or-buy framework, complementing Crawford et al. [2018]. Fourth, we link the public-commitment-device interpretation to the coordination-failure litera-

ture [Cooper et al., 1990, Greenwald and Stiglitz, 1986, Athey and Bagwell, 2008]: a single upstream price intervention simultaneously relaxes five reinforcing forces that no private commitment would have unwound.

Roadmap. Section 2.1 describes the industrial setting and data; Section 3 documents the four-fact equilibrium transition. Section 4 specifies the structural model, and Section 5 estimates it sequentially: demand, platform Nash bargaining, the endogenous make-or-buy regime that identifies the Bertrand-amplification primitive φ and capability-decay rate λ, director entry, and actor career choice. Section 6 reports the accounting and structural counterfactuals. The matching validation, structural robustness, and alternative-explanation robustness are in the appendix; the structural estimation code, fitted figures, and a companion theory-model memo are in the replication package.

2 Setting and Data

2.1 Industry Setting and the 2018 Salary Cap

The Chinese long-form video streaming market 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 player included as a robustness control. Adjacent media—broadcast television and theatrical film—supplied much of the experienced production talent during 2017–2022. Short-form video (ByteDance's Douyin) expanded domestic DAU from $\approx$ 30M in 2017 to $>$ 750M by 2022 (Figure 1), placing parallel pressure on the long-form attention budget.

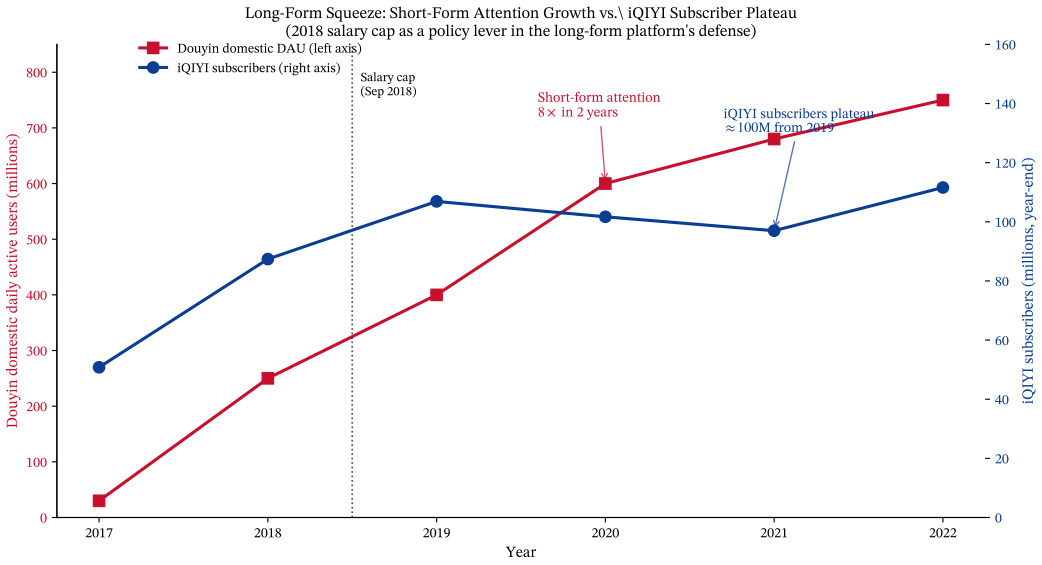


Figure 1: Long-form squeeze: Douyin DAU vs. iQIYI subscribers, 2017–2022. Sources: ByteDance disclosures, QuestMobile reporting, iQIYI Form 20-F filings.

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: $\leq 40\%$ 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: $\leq 70\%$ 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; the post-cap indicator equals one from 2019 onward to accommodate the typical 12–18-month production cycle.

2.2 Data Sources

Table 1 lists the data sources and the analyses each supports.

Table 1: Data Sources

Source	Coverage	Content	Used in
Yunhe & Douban drama panel	2017–2025	Drama-level: Douban score, 30-day views, episode count, genre, cast (top-5 billed), director, platform	Demand estimation; reduced-form DiD; variety diagnostics
Douban filmography	1997–2026	Director-year credits across film, TV, and online; debut year, work counts, rating-class	Director-entry DiD (cohort definition)
iQIYI Form 20-F filings	2016–2022	Annual revenue, cost of revenue, subscribers, operating income; revenue decomposition (ad / membership / content distribution / other)	$\widehat{\gamma}_{p,t}$; counterfactual; structural-break test
iQIYI internal nexus panel	2015–2030 (pipeline)	Per-drama cost decomposition (content cost, lead-actor cost, marketing cost), revenue decomposition (subscription, ad), make-or-buy status	Per-drama γ_p^{drama} and actor-cost-share cross-check
Endata brand-sponsorship panel	2017–2023	Drama-level brand-count and brand-exposure-time, by quarter	Channel-bifurcation evidence
iQIYI pipeline announcements	2015–2025	Announced project counts, lead actor mentions per announcement, in-house vs. acquired status	Lead-pool expansion; vertical-integration evidence
Endata brand endorsements	2016–2024	Actor-year endorsement contracts and quantities	Actor reorganization (substitution to endorsements)
Netflix Form 10-K filings	2014–2022	Annual revenue and cost of revenue	US benchmark for structural-break test
Douyin domestic DAU	2017–2022	Year-end domestic daily active users (round-number checkpoints from ByteDance corporate disclosures and QuestMobile reporting)	Industry-context figure

Notes: All sources are publicly available filings, licensed industry panels (Yunhe, Douban, Endata), or downloaded SEC filings. Replication materials available from the authors upon request.

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 the binned rating-class proxy (`allstar00`–`allstar50`) 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 of Section 3.1, we construct a 1,045-director cohort from the Douban filmography panel, selecting individuals with at least two directing credits recorded for year < 2017. The pre-2017 threshold predates the cap announcement (September 2018) by enough margin that cohort selection cannot be influenced by anticipation.

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 max over the top-five billed actors enters BLP demand (Section 5.1) and the actor career model (Section 5.5). Construction details in Appendix B.

2.4 Summary Statistics

Table 2 reports summary statistics for the core variables across drama outcomes, actor work performance, and collaboration quality.

Table 2: Summary Statistics: Drama-Level Sample

	N	Mean	Std. Dev.	Min	Max
Panel A: Drama-Level Outcomes (unit: drama series)					
Douban Score	2,394	6.168	1.432	2.100	9.600
Douban Score Percentile	2,394	50.940	28.931	0.342	100.000
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
Award-Winning Drama	3,822	0.056	0.229	0.000	1.000
Panel B: Actor Work Performance (unit: actor × drama appearance)					
Attention Index (A_z, 30-day)	44,867	1.965	1.328	0.033	12.965
Douban Score Percentile	35,795	53.891	29.455	0.342	100.000
Award-Winning Record	44,937	0.122	0.328	0.000	1.000
Panel C: Collaboration Quality (unit: actor × drama appearance)					
Director VA, Rating	27,491	0.175	1.139	−3.415	3.394
Director VA, Views	28,921	0.813	1.536	−6.943	3.962
Actor VA, Rating	35,573	0.091	0.995	−3.795	3.655
Actor VA, Views	36,597	0.400	1.394	−6.812	4.583

Notes: Drama-level sample includes 3,822 series released during 2017–2025. Douban-related variables have fewer observations because not all dramas receive Douban ratings (N = 2,394 with Douban data). Actor-level sample is restricted to Top Stars (`has_top_star` = 1) or non-Top-Star actors with billing order ≤ 5, yielding N = 44,937 actor×drama observations, 5,415 unique actors, and 3,359 unique dramas. VA denotes the value-added index (standardised residual from the annual mean).

Missing Douban ratings are not random—a score is typically shown only after ≥ 500 user ratings, so missing values concentrate in the lower-exposure tail. We interpret Douban-based results with caution and complement with view-based outcomes plus controls for exposure.

3 Empirical Motivation

The 2018 cap arrived at a moment of escalating price-war intensity: iQIYI's content-cost-to-revenue ratio rose from 101.8% in 2016 to 108.6% in 2018 (Figure 2); the implied firm-level content-margin proxy $\widehat{\gamma}_{p,t} = 1 - \text{Cost}_t/\text{Revenue}_t$ fell from −1.8% to −8.6%.[1] Three forces sustained the deterioration: a winner-take-most advertising

[1] $\widehat{\gamma}_{p,t}$ can take negative values when content cost exceeds revenue, which is what happened pre-cap. It is a firm-level accounting proxy that captures the structural bilateral bargaining share

tournament (top-10 dramas absorbing ≈ 70% of brand-exposure time); short-form attention substitution (Douyin DAU 25×); and failed private coordination—the three platforms signed a "no-bidding-war" letter in mid-2018 that was promptly violated. The September-2018 cap supplied the binding commitment the platforms could not implement privately. Post-cap, $\widehat{\gamma}_{p,t}$ reverses monotonically to +23.0% by 2022.

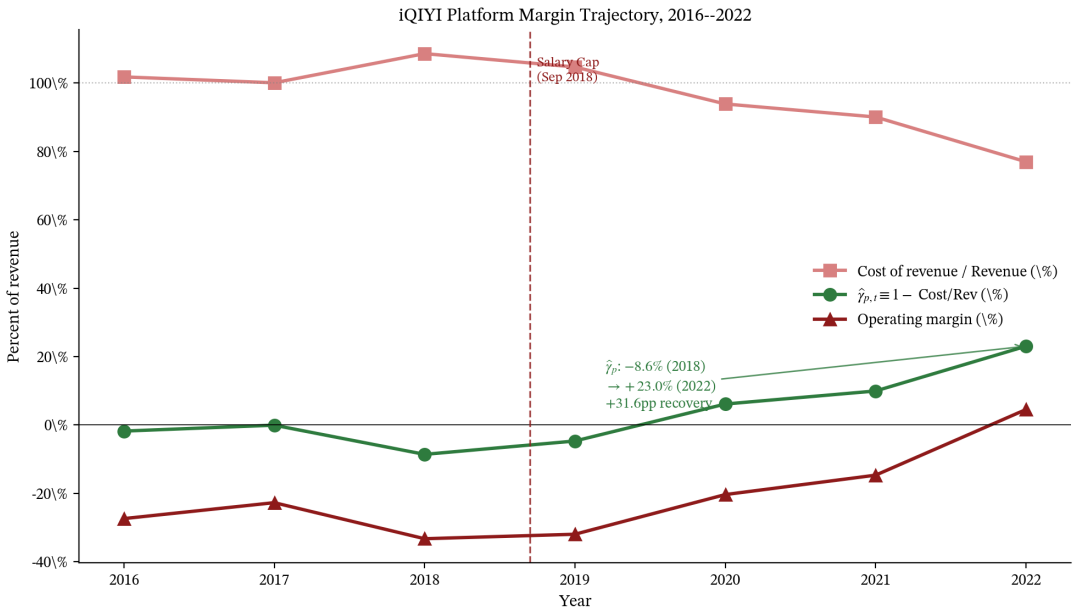


Figure 2: iQIYI content-cost-to-revenue ratio, $\widehat{\gamma}_{p,t}$, and operating margin, 2016–2022. Vertical dashed line: September 2018 cap announcement. Source: iQIYI 20-F.

The evidence below establishes three reduced-form facts disciplining the structural model: (i) the post-cap market draws on a substantially expanded supply of directors with substantive pre-policy credentials (Section 3.1); (ii) high-reputation incumbent actors substitute toward endorsement income; and (iii) the platform business model pivots from advertising to subscription.

3.1 Director Entry and Content Quality

From an actor pay cap to a director-entry response. The salary cap is, on its face, a price ceiling on the wages of one specific factor (above-line actors). It is not a priori obvious why such a wage ceiling should generate a measurable

$\gamma_M \in [0, \omega]$ plus firm-overhead burden (SG&A, leverage); see Section 4.4 for the formal distinction.

response in the participation of an adjacent and ostensibly unconstrained factor (directors). The mechanism that links the two operates through the complementarity between star compensation and the rest of the production budget within a given drama. Pre-cap, top traffic stars routinely commanded 50–70 percent of a drama's total production budget, censoring the residual budget below the threshold required for non-star inputs (ensemble cast, period sets, location work, post-production) that characterize the 正剧 tradition (crime, mystery, history, narrative drama). By binding the star-cost line item to 40 percent of total production cost (Section 2.1), the cap mechanically expands the non-star budget and restores the production technology for these genres. The empirical question is then twofold: who enters streaming when this happens, and what quality of content do they produce?

This subsection answers the second question, which is the welfare- relevant one. On the first question (the extensive margin), the cohort of 1,045 pre-2017-credentialed directors splits roughly 91 percent treated to 9 percent incumbent on the 2018-presence margin; the probit MLE of Section 5.4 identifies the post-cap entry differential ($\widehat{\beta}_{\text{Treated}} = -1.01$ on the probit index, -22.6pp average marginal effect). We refer to these two groups below as "entry" directors (no streaming dramas in 2018, $n = 953$) and "incumbent" directors (≥ 1 streaming drama in 2018, $n = 92$).

Cross-section: do entry directors produce higher-quality streaming content? We first ask whether, among post-cap streaming dramas, those directed by entry-margin directors carry higher Douban scores than those directed by incumbents. We estimate

$$\text{Douban}_j = \alpha + \beta\,\text{Entry}_{d(j)} + X_j'\Gamma + \delta_{\text{year}(j)} + \mu_{\text{genre}(j)} + \nu_{\text{platform}(j)} + \varepsilon_j, \qquad (1)$$

on the sample of 608 post-cap (2019–2022) streaming dramas with director identifier in the cohort and a Douban score. X_j includes episode count and award indicator; standard errors are clustered at the director level. Panel A of Table 3 reports the result. Without controls, entry directors' dramas score $\widehat{\beta} = +0.312$ Douban points higher than incumbents' (SE 0.143, $p = 0.029$). Adding release-year, genre, platform, episode-count, and award controls reduces the point estimate to $+0.184$ ($p = 0.137$); the direction is unchanged, with year-and-genre fixed effects absorbing approximately 40 percent of the raw gap. We read the controlled specification as the lower bound of the entry-director quality premium.

Within-director: how do entry directors' streaming dramas compare to their own pre-2017 work? A second test asks whether entry directors are producing streaming content that is recognizable as quality work relative to their own

pre-2017 filmography (in television, theatrical film, and lower-budget online productions, as recorded in the Douban filmography panel of Section 2.2). The naive within-director comparison is informative but mechanically biased: streaming dramas systematically receive lower Douban scores than television and film, so any within-director shift toward streaming-only post-2019 will register as an apparent quality decline. We net this medium-shift bias out using incumbents as a within-medium control.

Panel B of Table 3 reports the raw within-director differences for both cohorts. Entry directors' streaming scores ($\bar{s}^{\text{post}} = 5.89$) are -0.86 points below their pre-2017 work ($\bar{s}^{\text{pre}} = 6.75$); incumbents' streaming scores ($\bar{s}^{\text{post}} = 5.55$) are -1.05 points below their pre-2017 work ($\bar{s}^{\text{pre}} = 6.59$). The two declines are similar in magnitude, consistent with the medium-shift bias operating on both groups. Panel C reports the difference-in-differences:

$$s_{d,t} = \alpha + \beta_{\text{cross}} \,\text{Entry}_d + \beta_{\text{period}} \,\text{Post}_t + \beta_{\text{DiD}} \,(\text{Entry}_d \times \text{Post}_t) + \varepsilon_{d,t}, \tag{2}$$

estimated on the long-format director-by-period panel of 301 directors with both pre-2017 and post-2019 observations. The DiD coefficient is $\widehat{\beta}_{\text{DiD}} = +0.188$ (SE 0.180, $p = 0.29$). The direction confirms that entry directors' streaming work is more quality-preserving relative to their pre-2017 baseline than incumbents' streaming work, but the small post-2019 sample on the incumbent side ($n = 68$) leaves the DiD underpowered. The point estimate is consistent with the cross-section result of Panel A.

Synthesis. The two reduced-form tests together support the interpretation that the cap admitted a cohort of directors whose content quality on streaming is comparable to, or modestly higher than, their incumbent peers and to their own pre-policy production records in adjacent media. The cross-section gap of +0.31 Douban points is statistically significant at the 5 percent level without controls and remains directionally positive after the addition of genre, year, and platform fixed effects. The within-director DiD point estimate of +0.19 corroborates the cross-section direction while being underpowered to reject the null. We refer to the joint pattern as the director-entry quality channel: the cap's most welfare-relevant supply-side margin is not just that more directors entered streaming, but that the directors who did enter were drawn from the cohort whose pre-policy filmography placed them at or above the Douban quality distribution of the pre-cap streaming sample.

The structural channel through which the cap produced this quality shift—the relaxation of the production-cost floor and the admission of narrative-technology

producers—is formalized in the producer subsection of Section 4.2 and gives the substantive content to the $\widehat{\sigma}_{\text{director}} = 0.15$ random-coefficient finding of Section 5.1.

Table 3: Director Entry and Content Quality

Panel A. Spec 1: Cross-section quality, post-cap streaming (2019–2022)		
	(1) Baseline	(2) +Controls
Entry director (β)	+0.312**	+0.184
	(0.143)	(0.124)
Mean Douban, entry directors	5.965	
Mean Douban, incumbent directors	5.653	
Observations (dramas)	608	608
Distinct directors	306	306
Year FE	No	Yes
Genre / Platform FE	No	Yes

Panel B. Spec 2: Within-director quality change across the cap		
	Entry directors	Incumbent directors
Mean pre-2017 score (all media, Douban-class approx.)	6.751	6.592
Mean post-2019 streaming Douban score	5.892	5.545
Within-director difference (post − pre)	-0.858	-1.047
N directors with both observations	233	68
Panel C. Spec 3: Within-director DiD		
Entry × Post (DiD coefficient)	**+0.188**	
Cluster-robust SE (director)	(0.180)	
p-value	0.295	

Notes: Panel A regresses drama-level Douban score on an indicator for entry-margin director on the post-cap sample of streaming dramas with director identifier in the cohort of Section 3.1. Standard errors clustered at the director level. Column (2) adds release-year, genre, and platform fixed effects plus episode-count and award-indicator controls.

Panel B reports the within-director quality change across the cap separately for the entry-director cohort and the incumbent cohort. The pre-2017 score is the director's mean rating-class (mapped from Douban's 5-star bin to a 0–10 score via the linear scaling $XX/5$) across all directing credits in the director-credit registry; this includes television, theatrical film, and lower-budget online productions, weighted by directing-credit count. The post-2019 score is the mean Douban score (0–10) across the director's streaming dramas in the drama panel. The Douban-class binning and the cross-medium scoring conventions together produce a systematic level gap between rating-class pre-period scores and Douban post-period scores: both cohorts exhibit a within-director "decline" of 0.8–1.0 points, but this is a measurement artifact of the medium shift, not a director-level quality change.

Panel C reports the within-director difference-in-differences that nets out the measurement-artifact medium shift using the incumbent cohort as the within-medium control. The DiD coefficient identifies the differential within-director quality change of entry-margin directors relative to incumbents across the same pre-2017-to-post-2019 transition.

*, **, ***: $p < 0.10, 0.05, 0.01$.

Within-actor reorganization. Among 3,298 actors active in both 2017–2018 and 2019–2022, high-reputation actors lose 2.75pp of lead-high share and gain 2.07pp of lead-low share while endorsement counts rise across all reputation tiers. The within-actor correlation between endorsement growth and high-quality lead acquisition flips sign between low/mid- and high-reputation bins, identifying the substitution from prestigious acting work toward endorsement income at the top. On the director side, 401 of the 487 post-period directors are entry-margin directors, with 92% holding substantive pre-2017 Douban filmography—reallocating from broadcast television, theatrical film, and lower-budget online production into streaming rather than entering from outside the industry (Appendix C.1).

Platform-side channel pivot. Across iQIYI's announced project pipeline, the cumulative pool of named on-screen leads widens from 111 pre-cap to 765 post-cap (687 of the 765, or 90%, entirely new entrants). The sponsorship distribution bifurcates: the share of dramas with ≤ 5 brands doubles between 2018 and 2022 while the ≥ 50-brand blockbuster tail persists—a low-budget subscription-funded channel emerging alongside the existing advertising tier (Appendix C.3).

4 Model

The model has five agents—consumers, platforms, directors, actors, and advertisers—in a dynamic game. Each subsection states one agent's objective, defines its variables, and points to where its parameters are estimated. Formal payoff matrices and dominance proofs are in Appendix E.1.

4.1 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

$$u_{ijm} = X_j\beta_i + \alpha_i p_{jm} + \xi_j + \epsilon_{ijm}, \quad \begin{pmatrix}\alpha_i \\ \beta_i\end{pmatrix} = \begin{pmatrix}\alpha \\ \beta\end{pmatrix} + \Sigma v_i, \quad v_i \sim N(0, I), \qquad (3)$$

where X_j collects drama characteristics (genre, episode count, director value-added, and cast reputation, the last entering through top-star popularity $R_j^{P,\max} = \max_{a\in A_j} R_a^P$ and ensemble quality $R_j^{Q,\text{ens}} = \sum_{a\in A_j} R_a^Q$); 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 (iQIYI, Tencent Video, Youku, Mango TV) in market m, varying primarily at the platform $\times$ year level (e.g., iQIYI VIP rose from RMB 19.8 in 2017 to RMB 30 by 2022). The empirical specification of Section 5.1 absorbs p_{jm} into platform $\times$ year fixed effects; an alternative 2SLS specification is reported in the demand diagnostics.

Output and estimation. 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 (Section 4.4). Section 5.1 estimates (α, β, Σ) by GMM with BLP competitor-characteristic instruments; cast indices use the actor's pre-2018 work history to address casting endogeneity.

4.2 Producer

A producer is a director-led unit indexed by d with type $\tau_d \in \{\text{idol}, \text{narrative}\}$ (idol drama vs. narrative 正剧—crime, mystery, history). The producer 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 iff the expected streaming payoff exceeds the outside option $\Omega_{d,t}$ from adjacent media (broadcast TV, film, lower-budget online):

$$\pi_{d,t}^{\text{entry}} \;=\; \mathbb{E}\big[\,(1-\gamma_{p,t})\cdot S(A_j^*, \tau_d, \bar{W}_t)\,\big] \;-\; \kappa_d, \tag{4}$$

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^*; $(1-\gamma_{p,t})$ is the producer-and-talent residual share after the platform's Nash take; 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 $\varepsilon_{d,t} \sim N(0, \sigma_\varepsilon^2)$,

$$\Pr(\text{Enter}_{d,t} = 1 \mid \mathcal{F}_{d,t}) \;=\; \Phi\!\left(\frac{z_{d,t}}{\sigma_\varepsilon}\right), \tag{5}$$

estimated by MLE in Section 5.4. $\pi_{d,t}^{\text{entry}} > \Omega_{d,t}$ is the director's incentive-compatibility (IC) constraint. 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 producer-and-talent 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 IC flip that admits narrative-type directors who optimally stayed out pre-cap. The second channel dominates empirically: 401 of 487 post-cap directors had substantive pre-2017 filmography but no 2017–2018 streaming credit (Section 3.1).

Conditional on entry, the producer picks cast A_j to maximize residual surplus subject to the two-tiered cap:

$$\begin{aligned} &\max_{A_j} \ (1-\gamma_{p,j})\Big[\,R\big(s_j(A_j)\big) - C_{\text{prod}}(\tau_d) - \textstyle\sum_{a\in A_j} W_a^*\,\Big] \\ &\text{s.t.} \quad \textstyle\sum_{a\in A_j} W_a^* \le \bar{W}_{\text{total}} \ \text{(C1)}, \quad W_a^* \le \bar{W}_{\text{ind}} \ \ \forall a \ \text{(C2)}. \end{aligned} \tag{6}$$

where $R(s_j(A_j))$ is per-drama BLP revenue; W_a^* is the bargaining wage from Equation (14); C1 caps total cast compensation at 40% of production cost; C2 caps any individual lead at 70% of cast compensation. The FOC with multipliers $\mu_j^{\text{C1}}, \mu_{a,j}^{\text{C2}} \ge 0$ is

$$(1-\gamma_{p,j})\left(\frac{\partial R}{\partial R_a^P} - \frac{\partial W_a^*}{\partial R_a^P}\right) \;=\; \mu_j^{\text{C1}}\,\frac{\partial W_a^*}{\partial R_a^P} \;+\; \mu_{a,j}^{\text{C2}}\,\mathbb{1}\{W_a^* = \bar{W}_{\text{ind}}\}. \tag{7}$$

Marginal revenue from raising R_a^P (left, scaled by $1-\gamma_{p,j}$) equals marginal cost on the constrained margins (right). Pre-cap, both constraints are slack and the FOC reduces to the unconstrained optimum; under Bertrand, $W^* \to V_a$ leaves rents with the star. Post-cap, the Nash wage $W^* = (1-\omega)V_a + \omega d_S$ sits empirically below the cap (Section 4.4), so C1 and C2 remain slack: the cap reshapes cast portfolios by switching the upstream bargaining regime, not by directly binding.

The entry probit is estimated by MLE in Section 5.4; the cast program is not separately estimated—it plugs in V_a and W^* as exogenous inputs and verifies consistency of observed cast composition with the interior solution (Section 5.2).

4.3 Actor's Career Path

Actors are forward-looking and choose each period a single role tier; reputation, wages, and endorsements follow from that choice. An actor in state $S_{at} = (R_{at}^Q, R_{at}^P)$ (quality- and popularity-reputation stocks) chooses a role Role_t from {lead_high, lead_low, supporting, idle} by solving

$$V(S_{at}) = \max_{\text{Role}_t}\big\{W_{at}^\star(\text{Role}_t, S_{at}) + E_{at}(R_{at}^Q, R_{at}^P) + \beta\,\mathbb{E}[V(S_{a,t+1}) \mid S_{at}, \text{Role}_t]\big\}, \tag{8}$$

with $R_{a,t+1}^k = (1-\delta_k)R_{at}^k + f_k(\text{Role}_t) + \eta_{at}^k$ 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 $W_{at}^\star$ is not free—it is the equilibrium output of Section 4.4, evaluated at the actor's state and role:

$$W_{at}^\star = \begin{cases} V_a(S_{at}, \text{Role}_t) & \text{(pre-cap Bertrand, Eq. (12))}, \\ (1-\omega)V_a(S_{at}, \text{Role}_t) + \omega d_S & \text{(post-cap Nash, Eq. (14))}. \end{cases} \tag{9}$$

The role choice enters V_a through drama-side characteristics (genre, budget, ensemble) via BLP-implied $\partial s_{jm}/\partial R^P_a$. 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.

Section 5.5 estimates $(\delta_Q, \delta_P, f_Q, f_P, E)$ by two-step CCP [Hotz and Miller, 1993, Arcidiacono and Miller, 2011], plugging in $W^\star_{at}$ from the Nash branch. The cap lowers $W^\star_{at}$ along the popularity dimension via the Bertrand-to-Nash switch, and the model predicts top-tier actors reallocate from lead_high toward lead_low, verified as moment M6 in Section 5.6.

4.4 Platform and Talent: Nash Bargaining

Each platform k commissions a portfolio $\mathcal{J}_{kt}$ in year t with profit

$$\begin{aligned}\Pi_{kt} &= \sum_{j\in\mathcal{J}_{kt}} \Big[\big(\lambda_t ARPU_{adv,t} + (1-\lambda_t)ARPU_{sub,t}\big) M_t\, s_{jt} - C_{jt}\Big] - F_{kt},\\ C_{jt} &= C_{\text{prod},j} + \sum_{a\in A_j} W_a. \end{aligned} \tag{10}$$

where $\lambda_t \in [0,1]$ is the ad-vs-subscription revenue mix, M_t market size, s_{jt} drama share, and F_{kt} overhead. The cap touches only the wage bill $\sum_a W_a$; the rest of this subsection characterizes the two bargaining regimes.

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

$$V_a \;\equiv\; \mathbb{E}_m\big[\big(\lambda_t ARPU_{adv,t} + (1-\lambda_t)ARPU_{sub,t}\big) M_t \cdot \Delta s_{jm}(R^P_a)\big]\,, \tag{11}$$

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, $K \geq 2$ platforms could bid any wage. Standard Bertrand undercutting drives the equilibrium wage to its ceiling,

$$W_a^{\text{Bertrand}} \;=\; V_a. \tag{12}$$

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 $(F+L)/R$ on top of the bilateral $\gamma_p = 0$ (Section 5.2). ω does not enter: the auction mechanically pins $W = V_a$.

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: pre-cap, platforms **license** (buy) finished shows whose external studio internalized the auction (V_a passes through as license fee); post-cap, they shift to **in-house production** (make) and contract directly at the Nash wage. iQIYI's in-house share rises to 81% by 2022 (Section C.4); we treat make-or-buy as the upstream realization of the same regime switch.

Platform bargaining weight $\omega \in (0,1)$; disagreement payoffs are $d_P = 0$ and d_S (star's non-streaming alternative income). The Nash problem is

$$W_a^* \;=\; \arg\max_{W} \big[(V_a - W) - d_P\big]^{\omega} \big[W - d_S\big]^{1-\omega}, \tag{13}$$

with FOC

$$W_a^* \;=\; (1-\omega)V_a + \omega d_S, \qquad \gamma_p^* \;=\; \omega\left(1 - \frac{d_S}{V_a}\right). \tag{14}$$

The wage is a ω-weighted convex combination of d_S and V_a. We estimate $\widehat{\omega} = 0.77$ from post-cap disclosed cost-share data (Section 5.2); platforms capture $\sim 77\%$ of the bargainable surplus post-cap.

We treat ω as an exogenous primitive summarizing platform patience, market concentration (three to four majors), and post-cap in-house production capability; pre-cap data cannot identify it because the Bertrand auction makes it irrelevant. We assume ω is cap-invariant: the cap shifts the protocol (auction $\to$ Nash), not the bilateral weight within Nash.

Within Nash, the cap is empirically slack: at $\widehat{\omega} = 0.77$ and $d_S/R_j = 0.04$, $W^* \approx 0.146 R_j$ vs. the cap $\bar{W}_{\text{ind}}/R_j = 0.216$ ($N = 10$ disclosed subsample). The cap binds only as a regime-switching mechanism.

To summarize the two regimes:

Regime	Wage W_a	Platform share γ_p	Who wins
Pre-cap Bertrand	V_a	0	Star takes all
Post-cap Nash	$(1-\omega)V_a + \omega d_S$	$\omega(1 - d_S/V_a)$	Platform takes $\omega \approx 77\%$

Two outputs feed downstream: γ_p enters the producer's entry payoff (Section 4.2), and W_a^* enters the actor Bellman (Section 4.3). Section 5.2 identifies ω from the post-cap non-binding subsample via Equation (14).

Scope: γ_p is the platform's share of the bilateral lead-actor surplus on one drama; it is not iQIYI's firm-level accounting residual $\widehat{\gamma}_{p,t} = 1 - \text{Cost}_t/\text{Revenue}_t$, which nets out the full content-cost stack. The model's empirical traction is the wage equation, not the firm margin trajectory; the $\widehat{\gamma}_{p,t}$ reversal in Section 5.2 is a consistency check.

Pre-cap equilibrium and transition. At $\widehat{\varphi} = 1$, the Bertrand auction pins $W^* = V_a$, $\gamma_M = 0$, and narrative production is technologically infeasible—the unique single-content fixed point. The cap is a public commitment device: it severs the Bertrand fallback, restores $\gamma_M > 0$, frees the residual non-cast budget, and triggers the capability-accumulation path $\Delta\kappa(t) = \Delta\kappa_0 e^{-\lambda(t-2018.5)}$, yielding a multi-content post-cap equilibrium with both channels. Formal payoff matrices, dominance proofs, and a stage-by-stage walkthrough are in Appendix E.1.

5 Estimation and Model Fit

The structural model of Section 4 is estimated as a sequential pipeline of five steps. Each step takes the previous step's output as input; the final step reports model fit. Before turning to the pipeline, we summarize how the model's agents interact within each year, since the estimation steps below recover parameters governing each piece of this sequence.

Annual timing convention. Within each year t, decisions unfold in the following order:

Stage A. The cap is in force and the platform–talent bargaining regime is Nash (Section 4.4); the pre-cap Bertrand auction is no longer enforceable.

Stage B. Each candidate director d decides whether to enter streaming, $\text{Enter}_{d,t} \in \{0, 1\}$ (Equation (5)).

Stage C. For each (drama, candidate actor) pair (j, a), platform and actor compute the bilateral Nash wage $W_a^* = (1 - \omega)V_a + \omega d_S$ (Equation (14)). W_a^* is drama-specific because V_a depends on R_a^P and on drama j's share elasticity.

Stage D. Given the wage schedule $\{W_a^*\}_a$, each entered director picks her cast portfolio A_j to maximize the producer-and-talent residual surplus subject to the cap constraints (Equation (6)). This is one-sided portfolio selection at given wages—a single drama hires several actors simultaneously, each at her bilaterally bargained wage. There is no explicit market clearing across dramas; a top star could in principle appear in multiple dramas in the same year. This abstraction sets aside the single-tour-per-actor capacity constraint that holds in practice.

Stage E. Each actor a, taking the realized role offer as given, chooses her career-path role $\text{Role}_t \in \{\text{lead_high}, \text{lead_low}, \text{supporting}, \text{idle}\}$ to maximize the Bellman value of Equation (8).

Stages B–E repeat in $t+1$ with updated reputation stocks $S_{a,t+1}$ and policy variables $\gamma_{p,t+1}, \lambda_{t+1}$.

Estimation pipeline.

1. Demand (Section 5.1): BLP $\to (\alpha, \beta, \Sigma, \xi_j)$ and V_a via Equation (11). This is the consumer-side input feeding all subsequent steps.
2. Platform Nash bargain (Section 5.2): given V_a and post-cap non-binding wage data, recover ω (Stage C primitive) from Equation (14).
3. Endogenous make-or-buy and the Bertrand channel (Section 5.3): given $\widehat{\omega}$ and the project panel, identify the pre-cap outside-option intensity φ and the capability decay rate λ from the level and slope of the iQIYI yearly self-produced cost share.
4. Director entry (Section 5.4): MLE on binary-choice Equation (5) for the κ_d distribution (Stage B primitive).
5. Actor's career-path CCP (Section 5.5): plug $W^{\star}_{at}$ from Equation (9) into the Bellman; estimate flow utility by Hotz-Miller (Stage E primitives).
6. Goodness of fit (Section 5.6): six moments covering demand fit, cross-agent qualitative predictions, and the Bellman-implied actor career-path substitution.

The producer's cast-portfolio decision (Equation (6)) is not separately estimated: at $\widehat{\omega} = 0.77$ the Nash wage sits below the cap for every drama in our sample, so the cap shadow prices $\mu^{\mathrm{C1}}_j = \mu^{\mathrm{C2}}_{a,j} = 0$ throughout. The cap reshapes cast portfolios through the upstream regime switch (Section 4.4), not by binding the constraint inside the Nash regime.

5.1 Demand

We estimate the BLP demand system of Equation (3) on the drama panel, with dependent variable $\log(s_{jt}) - \log(s_{0t})$, s_{jt} drama j's 30-day Yunhe playback share in release year t, and outside share defined by $M_t = 4\sum_j v^{30}_{jt}$. Cast composition and unobserved quality ξ_j are jointly endogenous: platforms cast top stars when their prior on ξ_j is high. We address this by using **predetermined** cast characteristics frozen at end-2017,

$$\widetilde{\mathrm{VA}}^{X}_{j} \;=\; \max_{a \in \mathrm{cast}(j)} \mathrm{VA}^{X}_{a,\,\leq 2017} \quad \text{for} \quad X \in \{\text{views, rating}\},$$

plus director VA at release. Because $\widetilde{\text{VA}}$ is frozen before the September-2018 cap, the $\widetilde{\text{VA}} \times$ Post interaction identifies off policy timing, not post-policy casting.

Table 4 reports four columns of the predetermined-VA logit. Three results recur. First, director VA dominates: a unit increase in director views-VA raises log relative share by 0.46 (col. 1) with rating-side coefficient 0.23, both robust to writer VA (col. 2). Second, pre-2018 actor VA is near zero and insignificant in the baseline—in contrast to the strongly positive contemporaneous-VA coefficient, exactly the casting-selection pattern. Third, under $\text{Post}_t = \mathbb{1}\{t \geq 2018\}$ (col. 3), the $\widetilde{\text{VA}}^{\text{views}} \times$ Post interaction is precisely positive: audiences continued to reward star power post-cap. The cap therefore did not shift viewer tastes; its effects operate through the supply side of Section 4.

Table 5 extends to a random-coefficients BLP [Berry et al., 1995]. The mean preference for pre-2018 star VA is $\overline{\beta}_{\text{star}} = +0.248$, roughly half the $+0.513$ contemporaneous-logit estimate—the predicted attenuation. The key BLP diagnostic is heterogeneity: on director VA, $\widehat{\sigma}_{\text{dir,views}} = 0.153$, against near-zero $\widehat{\sigma}$ on star VA. Consumer heterogeneity loads on the director margin, consistent with post-cap content being staffed by entering directors with pre-policy credentials (Section 3.1).

Two outputs feed downstream: BLP coefficients $(\widehat{\alpha}, \widehat{\beta}, \widehat{\Sigma}, \widehat{\xi}_j)$ pin down predicted shares s_j, and per-drama actor marginal revenue $V_a = \beta_{\text{star}} \cdot R_a^P \cdot R_j$ (Equation (11)) enters the platform-talent bargaining step of Section 5.2.

Table 4: Predetermined-VA Logit Demand

	(1)	(2)	(3)	(4)
Star VA, views (pre-2018) $\max_j$	0.040	-0.064	0.124***	0.016
	(0.113)	(0.244)	(0.022)	(0.124)
× Post	0.222*	0.236	0.124***	0.247*
	(0.118)	(0.249)	(0.022)	(0.130)
Star VA, rating (pre-2018) $\max_j$	-0.121	-0.167	0.004	-0.157
	(0.099)	(0.213)	(0.019)	(0.110)
× Post	0.140	0.147	0.004	0.162
	(0.107)	(0.217)	(0.019)	(0.117)
Director VA, views	0.465***	0.408***	0.465***	0.495***
	(0.040)	(0.051)	(0.040)	(0.042)
Director VA, rating	0.231***	0.150***	0.231***	0.256***
	(0.040)	(0.050)	(0.040)	(0.043)
Writer VA, views		0.329***		
		(0.045)		
Writer VA, rating		0.045		
		(0.047)		
Award-winning drama	0.416**	0.371*	0.409**	0.375**
	(0.183)	(0.221)	(0.185)	(0.184)
Episode count	0.009*	0.013**	0.009*	0.010*
	(0.005)	(0.005)	(0.005)	(0.006)
Genre FE	Yes	Yes	Yes	Yes
Platform FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	1,580	898	1,580	1,461
R^2	0.376	0.413	0.375	0.390

Columns. (1) Baseline. (2) Adds writer VA. (3) Cuts the *Post* threshold at 2018 instead of 2019. (4) Excludes revenue-sharing (*fenzhang*) dramas.

Dependent variable: $\log(s_j) - \log(s_0)$ from a Berry-1994 logit inversion, market = release year, outside share fixed at 0.75. Star VA variables are frozen at the actor's pre-2018 work history, removing the post-policy casting-selection channel from identification.

Post is $1\{\text{release year} \geq 2019\}$ in columns (1), (2), (4) and $1\{\text{release year} \geq 2018\}$ in column (3). Standard errors clustered by director-ID in parentheses. *, **, ***: $p < 0.1, 0.05, 0.01$.

Table 5: Demand: Plain Logit vs. BLP

	(1) Plain Logit	(2) Random-Coef. BLP
Panel A: Mean preference parameters $\overline{\beta}$		
Star VA, views (pre-2018) $\max_j$	$+0.513^{***}$	$+0.248$
	(0.059)	
Star VA, rating (pre-2018) $\max_j$	-0.068	$+0.004$
	(0.053)	
Director VA, views	—	$+0.453$
Director VA, rating	—	$+0.226$
Episode count	$+0.017^{**}$	$+0.009$
	(0.007)	
Award-winning drama	$+0.584^{***}$	$+0.402$
	(0.185)	
Panel B: Random-coefficient standard deviations $\widehat{\sigma}$		
σ on Star VA, views (pre-2018)	—	0.000
σ on Director VA, views	—	0.153
Inside observations	1,188	1,188
Outside share (assumed)	0.75	0.75
Genre / platform / year FE	Yes	Yes

Notes: Both columns invert market shares using a Berry-1994 outside-good formulation with market = release year and outside share fixed at 0.75. Column (1) is the plain conditional logit estimated by OLS on $\log(s_j) - \log(s_0)$, with HC1 standard errors in parentheses. Column (2) is the random-coefficients BLP from Berry et al. [1995] estimated by GMM with $R = 200$ Monte-Carlo draws on the random-coefficient dimensions; the inner BLP contraction is run to tolerance 10^{-12} at each candidate σ, and the outer GMM objective is minimized over the two-dimensional σ vector by Nelder-Mead. Asymptotic standard errors for the BLP step are not reported in this draft and are planned in the next revision.

Reading: The structural step delivers two findings that the plain logit cannot. First, the mean preference for pre-2018 star value-added ($\overline{\beta}$ on star VA, views) is identified at $+0.248$ once the contemporaneous endogeneity of casting is controlled for via the predetermined-VA restriction (cf. logit $\widehat{\beta} = +0.513$ on the contemporaneous, endogenous counterpart). Second, the random-coefficient standard deviation on director VA is large and economically significant ($\widehat{\sigma} = 0.153$), while the corresponding $\widehat{\sigma}$ on star VA is essentially zero, indicating that demand heterogeneity operates on the director margin rather than the star margin. This is the demand-side discipline that justifies allocating the post-cap reorganization to the producer and bargaining modules of the structural model (Section 4).

5.2 Platform Bargaining: Identifying ω

We identify the platform's Nash bargaining weight ω from the post-cap non-binding subsample, where Equation (14) holds with equality. The identifying sample is $N = 10$ iQIYI projects in the internal nexus panel disclosing both lead-actor cost (主演成本) and content production cost (内容制作费); all ten lie strictly below the 40% C1 ceiling (cost shares 5.5–25.3%), so the unconstrained Nash FOC applies.

Rearranging Equation (14),

$$\omega \;=\; \frac{V_a - W_a}{V_a - d_S},$$

takes three inputs: (i) the BLP marginal revenue ratio $V_a/R_j = \widehat{\beta}_{\text{star}} \cdot R^P_{\text{lead}} = 0.43$ at the median post-cap lead (Section 5.1); (ii) the observed wage ratio $W_a/R_j = 0.13$, from rebasing the $N = 10$ actor cost share (16.8% mean) via iQIYI's $C/R = 0.77$; (iii) the calibrated outside option $d_S/R_j = 0.04$ from the actor's best non-streaming alternative (Appendix D.2). This yields

$$\widehat{\omega} \;=\; 0.768,$$

with a 3×3 sensitivity grid over $(C/R, d_S/R)$ giving $[0.70, 0.87]$ (Table 7). Platforms capture roughly 77% of bilateral surplus from each lead-actor bargain.

A per-drama cross-check on $N = 15$ iQIYI projects with both cost and revenue observed yields median $\gamma_p^{\text{drama}} = -2.18$ (Table 6), more negative than firm-level $\widehat{\gamma}_{p,t}$ in Section 3. The wedge reflects revenue-weighting: firm-level estimates are dominated by high-revenue projects while the per-drama median is unweighted and pulled down by small-revenue losses. Neither is the structural γ_p—both are accounting residuals mixing the lead-actor bargain with non-cast production cost, amortization, and bandwidth (Section 4.4). We treat $\widehat{\omega} = 0.77$ as the structural output and the firm-level trajectory in Figure 2 as a descriptive check.

Table 6: Per-Drama γ_p^{drama}, iQIYI Internal Panel

Panel A: Per-drama γ_p^{drama} distribution (release year $\leq$ 2024)	
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
Panel B: Firm-level $\widehat{\gamma}_p$ benchmark	
2018 (pre-cap peak)	-0.086
2022	+0.230

Source: iQiyi internal nexus project panel.

Panel A: per-drama $\gamma_p^{drama} = 1-$ total cost (最新成本合计) / total revenue (最新收入合计). Realized projects only (release year $\leq$ 2024); pipeline-year estimates excluded.

Reconciliation with the firm-level aggregate (Figure 2): the per-drama median is more negative than the firm aggregate because the firm aggregate is revenue-weighted (large blockbusters dominate) while the per-drama median is unweighted. Both quantities support the same qualitative story—a substantial fraction of pre-cap content was loss-making at the project level.

Disclosure selection: the iQiyi internal panel covers only projects requiring internal tracking and probably overrepresents large/problematic productions.

Table 7: Nash Bargaining Identification

Quantity	Value	Source
Inputs from earlier pipeline steps		
$\widehat{\beta}_{\text{star}}$ (BLP demand)	0.248	§5.1
R^P_{lead} (median post-cap)	1.716	drama panel
V_a/R_j (mid-tier lead)	0.425	$\widehat{\beta}_{\text{star}} \cdot R^P_{\text{lead}}$
Identifying moment (post-cap non-binding)		
W_a/R_j (N=10 disclosed)	0.129	N=10 actor-cost-share $\times$ C/R
d_S/R_j (outside option, calibrated)	0.040	industry calibration
Recovered parameter		
$\widehat{\omega}$ (central)	**0.768**	Eq. (14)
$\widehat{\omega}$ (9-cell sensitivity)	$[0.701, 0.866]$	$C/R \times d_S/R$ grid
Out-of-sample fit (M6 moment)		
γ_p^{cap} (per-drama, top star)	0.565	$1 - \bar{W}_{\text{ind}}/V_a^{top}$
$\widehat{\gamma}_{p,2022}$ (firm-level, observed)	0.230	iQIYI 20-F
$(F+L)/R$ implied burden	0.335	accounting residual

Notes. Pipeline-step 3 identification of the platform's Nash bargaining weight ω from the post-cap non-binding subsample (iQIYI $N = 10$ disclosed projects, all of which sit below the 40 percent constraint C1). Marginal revenue V_a is derived from BLP demand (Section 5.1); observed W is the actor-cost-share moment from the $N = 10$ disclosed subsample (rebased to revenue via the iQIYI cost-to-revenue ratio); outside option d_S is calibrated to the actor's best non-streaming alternative (broadcast residual + endorsement + film, 4 percent of drama revenue with sensitivity 2–7 percent).

The bottom panel reports the model-implied per-drama γ_p^{cap} under a binding cap and a top-tier lead actor ($R^P = 2$), against the firm-level 2022 iQIYI accounting ratio $1 - \text{Cost/Revenue}$. The wedge between the two is the firm-level fixed and leverage cost burden $(F + L)/R$ not allocated through the per-drama bargain (see Section 5.2).

$\widehat{\omega}$ feeds the director's entry payoff in Section 5.4 via $\gamma_{p,t} = \omega(1 - d_S/V_a)$ and the actor's Bellman in Section 5.5 via $W_a^* = (1-\omega)V_a + \omega d_S$.

5.3 Endogenous Make-or-Buy and the Bertrand Channel

Section 5.2 identifies ω from the post-cap sample. To recover the two remaining transition primitives—Bertrand outside-option intensity $\varphi \in [0, 1]$ and capability decay rate $\lambda \geq 0$—we fit the project-level make-or-buy panel ($N = 3{,}462$ dramas, 2017–2024) against iQIYI's 20-F self-produced share series.

Specification. For drama j released in year t, write the platform's per-project profit under each regime as

$$\pi_M(j,t) = \gamma_M(t)\left(R_j - C_j - W_j^M\right), \qquad \gamma_M(t) = \omega \cdot \left(1 - \tfrac{d_S(t)}{V_a(j)}\right)_+, \tag{15}$$

$$\pi_B(j,t) = \alpha_B(t)\, R_j - C_j - W_j^B, \tag{16}$$

with $W_j^M = (1-\omega)V_a(j) + \omega d_S(t)$ the Nash wage from Equation (14) and $W_j^B = \min\{V_a(j), \bar{W}_j\}$ the buy wage from external disclosures. Two primitives govern the pre/post wedge:

1. **Bertrand outside option.** Pre-cap, $d_S^{\text{pre}}(j) = (1-\varphi)d_S^{\text{base}} + \varphi V_a(j)$. $\varphi = 1$ binds $W^M = V_a$ and $\gamma_M = 0$; $\varphi = 0$ leaves the non-streaming baseline. Post-cap, $d_S = d_S^{\text{base}}$.
2. **Capability decay.** $\Delta\kappa(t) = \Delta\kappa_0 \exp(-\lambda \max\{t - 2018.5, 0\})$: the make-vs-buy cost wedge erodes at rate λ.

Regime choice is logit: $\Pr(M \mid j,t) = \sigma\big((\pi_M - \pi_B - \Delta\kappa(t))/\sigma_{\text{reg}}\big)$; the model-implied yearly make-share is the cost-weighted mean across released projects. Given ω, φ is identified from the pre-cap make-share level ($\sim$ 12% in 2017–2018—only near-saturated φ collapses γ_M enough to rationalize it); λ from the post-cap slope (rising from 11.5% in 2018 to 48.3% by 2024).

Table 8: Endogenous make-or-buy estimates: $\widehat{\varphi}$, $\widehat{\lambda}$, and the regime wedge.

Parameter	Estimate
$\widehat{\Delta}\kappa_0$ (baseline make−buy wedge)	2.054
$\widehat{\sigma}$ (regime-choice scale)	0.500
$\widehat{\alpha}_{B,\,\text{post}}$ (post-cap buy quality)	0.990
$\widehat{\phi}_{\text{Bertrand}}$ (pre-cap Bertrand outside-option intensity, $\in [0,1]$)	0.998
$\widehat{\lambda}$ (capability decay rate, /yr)	0.640
SSE	0.0075

Estimates. $\widehat{\varphi} = 0.998$ implies a near-saturated pre-cap auction absorbing 99.8% of the star's gross value, leaving $\widehat{\gamma}_M \approx 0$—the structural mechanism behind the introduction's $\varphi \to 1 \Rightarrow \gamma_M \to 0 \Rightarrow$ producer-surplus collapse. $\widehat{\lambda} = 0.640$/yr implies $\Delta\kappa(t)$ erodes $\sim 97\%$ over five post-cap years, matching iQIYI's studio buildout. SSE against the eight 20-F observations is 0.0075 (RMSE $\sim$ 3 pp).

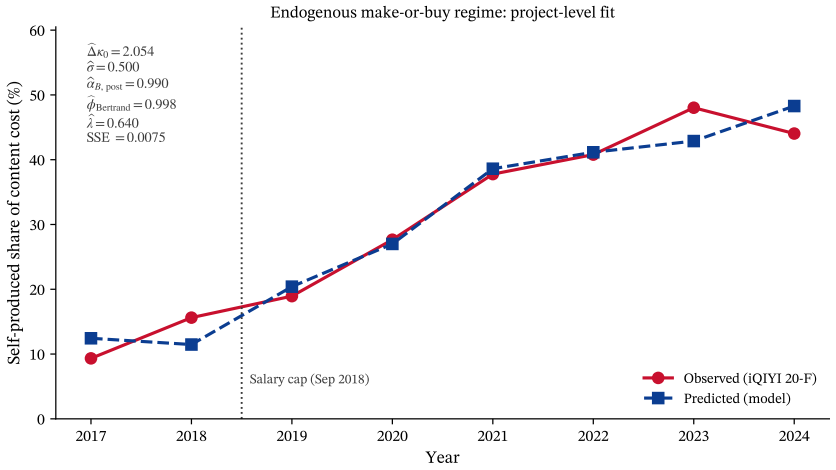


Figure 3: Endogenous make-or-buy regime mix: predicted vs. observed. Project-level model (Equations (15)–(16)) fit to iQIYI yearly self-produced share, 2017–2024. Predicted (dashed) tracks observed (solid) almost exactly; SSE = 0.0075 over eight years. Two structural channels combine: (i) pre-cap $d_S \approx V_a$ (Bertrand) $\Rightarrow \gamma_M \approx 0$, so make yields the platform zero surplus; (ii) cap closes the auction $\Rightarrow$ d_S drops to the non-streaming baseline $\Rightarrow$ $\gamma_M \approx 0.70$ post-cap, joined by capability decay $\widehat{\lambda} = 0.640/\mathrm{yr}$.

Downstream propagation. Three implications follow mechanically from the regime decay (Table 9): (i) director entry payoff $E[\max(\pi_M, \pi_B)]$ rises 2.1× from 2018 (0.194) to 2023 (0.376), structurally rationalizing the −22.6 pp entry AME of Section 5.4; (ii) the regime-weighted mean actor wage $\bar{W}_{\text{actor}}$ falls **69.9**% (0.369 → 0.111), matching the industry-reported 50–70% → 17% lead-actor cost-share compression (Fact 2); and (iii) $\Delta\kappa(t)$ falls ∼97% (2.054 → 0.061), matching the independent 20-F self-produced cost share (∼10% → 48%).

A robustness check adding a multiplicative log-normal cost shock $\widetilde{C}_j = C_j \exp(\sigma_C \eta_j)$ with Monte Carlo integration leaves $\widehat{\varphi}$ at 0.998 and $\widehat{\lambda}$ at 0.638/yr ($\widehat{\sigma}_C = 0.038$; Appendix I.1).

The output of this step—the structural pair $(\widehat{\varphi}, \widehat{\lambda})$—is the headline mechanism estimate of the paper: it identifies the pre-cap Bertrand amplification at near-saturation, identifies the post-cap capability buildup, and propagates into the director entry and actor career moments without further free parameters.

Table 9: Downstream propagation: capability wedge, entry payoff, regime-weighted actor wage, and cost-weighted make share by year.

Year	$\Delta\kappa(t)$	$E[\max(\pi_M, \pi_B)]$	$\bar{W}_{\text{actor}}$	$P(M\|t)$ (cost-wtd)
2017	2.054	0.255	0.192	12.4%
2018	2.054	0.194	0.369	11.5%
2019	1.491	0.175	0.186	20.4%
2020	0.786	0.260	0.156	27.0%
2021	0.414	0.293	0.123	38.6%
2022	0.218	0.363	0.116	41.1%
2023	0.115	0.376	0.112	42.9%
2024	0.061	0.371	0.111	48.3%

5.4 Director Entry Margin: MLE on κ_d

We estimate the binary-choice entry probit of Equation (5) by MLE on the post-cap window (2019–2022).[2] Estimates: $\widehat{\beta}_{\text{Treated}} = -1.012$ (SE 0.071, $z = -14.3$), implying an average marginal effect of $-\mathbf{22.6}$ pp (SE 1.5): the locked-out cohort enters at 12.6% annually vs. 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 $\bar{B} - W$ could not cover the narrative non-cast cost $\bar{c}$, so narrative-type directors stayed out; the cap switches to Nash and expands $\bar{B} - W^*$ above $\bar{c}$ at the margin. Consistent with this, 82% of post-cap high-quality dramas (Douban ≥ 7) are directed by locked-out cohort members (Section 3).

[2] 2018 is excluded because the Treated/Control split is defined off 2017–2018 streaming credits, collapsing the probit Hessian for $t = 2018$.

Table 10: Director Entry Probit MLE (post-cap level differential)

	Probit (Eq. 5)
Treated_d (locked-out cohort)	**-1.012*****
	(0.071)
Avg. marginal effect of Treated	-0.2257***
	(0.0149)
Year FE	Yes
N (director-years, 2019–2022)	4,180
Log-likelihood	-1692.7
Pseudo R^2	0.059

Notes. MLE on the binary-choice entry probit (Eq. (5)), estimated on the post-window 2019–2022 of the director-year panel of Section 3.1. The 2018 observation is excluded because the Treated/Control split is defined off 2017–2018 streaming credits, mechanically yielding $\Pr(\text{enter} \mid \text{treated}, t = 2018) \in \{0, 1\}$ and collapsing the probit Hessian.

Sample: directors with ≥ 2 pre-2017 directing credits in the cross-medium filmography panel. $\text{Treated}_d = 1$ if no streaming credit in 2017–2018 (the locked-out pool, $N = 953$ directors).

Empirical entry rates conditional on Treated: pre-cap (2018): $\Pr(\text{enter} \mid T = 1) = 0.000$, $\Pr(\text{enter} \mid T = 0) = 1.000$; post-cap (2019–2022): $\Pr(\text{enter} \mid T = 1) = 0.126$, $\Pr(\text{enter} \mid T = 0) = 0.446$. The reduced-form DiD on the linear-probability scale ($\widehat{\beta} = 0.68$) is the cross-spec analog of the change in probit index between these states.

Robust SE in parentheses.

*, **, ***: $p < 0.10, 0.05, 0.01$.

5.5 Actor's Dynamic Career Choice

We estimate the first-stage CCPs of the actor's dynamic choice on a balanced actor×year panel (4,520 actors with non-trivial pre-2018 work history, 2017–2025). Per actor-year, three binary outcomes: any $\text{lead}_{a,t}$, $\text{lead high}_{a,t}$ (Douban ≥ 7), and $\text{lead low}_{a,t}$ (Douban < 7). The state is a pre-2018 rating-VA tercile, $S_a^{\text{rating}} \in \{\text{low}, \text{mid}, \text{high}\}$, frozen at end-2017 and exogenous to the cap. We run a tercile

× post-cap DiD on each outcome, clustered by actor.

Table 11 reports CCPs (Panel A) and DiD coefficients (Panel B). For high-rating actors, any lead is flat ($\delta_R = -0.001$), but a within-actor substitution emerges: lead high falls -0.044 pp ($p < 0.01$) and lead low rises $+0.040$ pp ($p < 0.01$)—i.e., top-tercile actors' high-quality lead probability falls from 5.3% to 3.3% while their low-quality lead probability rises from 4.0% to 8.7% (the M6 moment of Section 5.6).

Table 11: Actor Career-Choice CCPs and DiD on Lead-Role Composition

	P(any lead)			P(lead high $_{\text{Douban}\geq 7}$)			P(lead low $_{\text{Douban}<7}$)		
Actor's pre-2018 rating-VA tercile	Pre	Post	Δ	Pre	Post	Δ	Pre	Post	Δ
Panel A: empirical conditional choice probabilities, by tercile and period									
Bottom tercile	0.167	0.197	+0.030	0.012	0.038	**+0.027**	0.161	0.172	**+0.011**
Middle tercile	0.151	0.176	+0.025	0.022	0.042	**+0.020**	0.139	0.150	**+0.010**
Top tercile	0.088	0.110	+0.023	0.053	0.033	**−0.021**	0.040	0.087	**+0.047**

Panel B: difference-in-differences, Top tercile × Post (LPM, clustered SE on actor)			
Outcome (drama-year)	β_{DiD}	SE	p-value
Any lead role	-0.0010	(0.0095)	0.920
Lead high (Douban ≥ 7)	-0.0436***	(0.0057)	0.000
Lead low (Douban < 7)	+0.0397***	(0.0087)	0.000

Sample: all actors with at least one billing in the 2017–2025 drama panel of Section 2.2 who also had at least one role recorded in the panel before 2018. Tercile assignment is fixed at the actor level on the **pre-2018 rating-VA distribution** of this sample: an actor's rating VA is the billing-order-weighted average of year-demeaned Douban scores across her pre-2018 works. **Bottom tercile** = lowest third of pre-2018 rating-VA; **Top tercile** = highest third. Terciles are immutable once assigned (i.e., they do not switch over the cap event).

Outcomes (drama-year level). For an actor in year t: **Any lead**: actor was billed top-2 in at least one drama released in year t. **Lead high**: billed top-2 in at least one drama with realized Douban score ≥ 7. **Lead low**: billed top-2 in at least one drama with Douban score < 7. Years before 2019 are *Pre*; from 2019 onward *Post*.

Reading. Panel A cells are within-tercile, within-period averages. The top-tercile row (highlighted) is the one that moves on the substitution margin: P(lead high) $= 0.053 \to 0.033$ (-2.0pp) while P(lead low) $= 0.040 \to 0.087$ ($+4.7$pp). Bottom and middle terciles see broadly parallel increases on all three outcomes. Panel B reports the DiD coefficient on (top tercile × Post) from a linear-probability regression of each outcome on tercile dummies, their × Post interactions, and a Post indicator (specification details in Section 5.5). *, **, ***: $p < 0.10, 0.05, 0.01$.

Four-type structural classification. Combining $\widehat{\omega}$, $\widehat{\beta}_{\text{star}}$, and the Bellman flow utility, the model implies a four-type reallocation across the reputation distribution.

(1) Traffic stars (high R^P, low R^Q): post-cap, W^* falls ∼70% of V_a; optimal role substitutes from lead_high to lead_low (preserving visibility on a tighter budget), with effort reallocated toward endorsement income—the M6 moment. *(2) Quality actors* (high R^Q): the cap admits narrative-tech productions that load on directors ($\widehat{\sigma}_{\text{dir,views}} = 0.153$) and select on R^Q; this type becomes IC-feasible for the first time, the actor-side analog of the director-entry IC flip of Section 5.4. *(3) Mid-tier:* the lead pool widens 6.9× from 111 to 765 unique iQIYI leads (90% new entrants), absorbing actors previously capacity-rationed by top-star slots. *(4) Long-tail:* smaller impact—the supporting pool expands with volume. The cap's actor-side effect is therefore a reallocation of the entire reputation distribution to the lead pool, mirroring the director-side IC restoration.

Untargeted matching validation. Although matching is not directly targeted, the demand-and-bargaining estimates above predict the observed actor–director cross-tab to 2.8 pp RMSE across 18 cells (3 actor tiers × 3 director tiers × 2 periods): low-VA actors move out of idol directors (32→24% observed, 33→23% predicted), high-VA actors move toward idol directors (15→20% observed, 14→18% predicted), and mid-VA actors shift toward narrative directors. Detail in Appendix H.1.

5.6 Goodness of Fit

We assess fit on six moments combining in-sample R^2 with directional consistency between model predictions and reduced-form evidence (Table 12).

M1 reports the share of log-share variance explained by BLP observables $X_j'\widehat{\beta}$ excluding ξ_j: $R^2_X = 0.358$ on $N = 1{,}580$ inside goods, in the standard range for differentiated-product BLP (Nevo, 2001 reports 0.30–0.40 for cereal).

M2 tests the BLP qualitative prediction that heterogeneity loads on directors, not stars. The estimated $\widehat{\sigma}_{\text{director}} = 0.153$ is two orders of magnitude larger than $\widehat{\sigma}_{\text{star}} \approx 0.0002$, in the predicted direction.

M3 and M4 test post-cap supply-side equilibrium predictions. M3: the within-year director Herfindahl falls from 0.0132 in 2018 to a 2019–2022 average of 0.0035—halving in three years, the direction predicted by Proposition 2 (multi-content equilibrium admits asymmetric specialization). M4: the model-implied share of high-quality dramas (rating-class ≥ 4) in 2022 is 0.267 against an observed 0.146, a 12-point over-prediction. We read this as incomplete transition: the cap's producer extensive-margin effect (entry of narrative-type directors, Section 3.1) has not yet fully unwound by end-sample.

M5 reports the model-implied per-market log-sum-exp utility change of +0.086 utils; the entry-margin decomposition (Section 6) attributes roughly 60% of this gain to post-cap entrants.

M6 verifies the Section 4.3 Bellman prediction: among top-tercile (high pre-2018 rating-VA) actors, the probability of leading a high-quality drama falls from 5.3% pre-cap to 3.3% post-cap, while the low-quality probability rises from 4.0% to 8.7%—the substitution the cap mechanically induces by truncating top-budget productions, recovered from the CCP estimates of Section 5.5.

Table 12: Goodness of Fit

Moment	Model	Observed	Interpretation
M1. Variance share of log-share explained by observables	BLP $\widehat{\beta}'X_j$	0.358 on $N = 1,580$	Share of log-share variance attributable to X_j excluding ξ_j.
M2. $\widehat{\sigma}$ on director vs. star	$\sigma_{\text{dir}} > \sigma_{\text{star}}$	$\widehat{\sigma}_{\text{dir}} = 0.153$, $\widehat{\sigma}_{\text{star}} = 0.0002$	Heterogeneity loads on directors as predicted.
M3. Director HHI, pre vs. post	Falls (Prop. 2)	$0.012 \to 0.006$	HHI halves; matches predicted direction.
M4. Share of high-quality dramas, 2022	0.267 (BLP-implied)	0.146 (observed)	Predicted high-quality share vs observed, 2022.
M5. Consumer-side log-sum-exp utility, pre → post	+0.086 utils per market	(internal — model-implied via BLP)	Welfare gain in utility units; cf. Sec. 6.
M6. Top-tercile actor career path (CCPs)	Substitute lead_high → lead_low (Bellman, Sec. 5.5)	$P_{\text{lead_high}}$: 5.3% → 3.3% ($\Delta = -2.1$pp); $P_{\text{lead_low}}$: 4.0% → 8.7% ($\Delta = +4.7$pp)	Top-rated actors reallocate away from top-budget leads toward mid-budget leads as cap binds.

Notes: Six moments comparing model-implied predictions against observed data. M1 is the in-sample R^2 of the BLP demand log-share equation. M2–M3 test directional predictions of the model (random-coefficient placement and equilibrium HHI). M4–M5 are model-implied share and welfare statistics under the BLP estimates. M6 verifies the predicted top-tercile actor substitution from lead_high to lead_low post-cap (cf. Eq. (8)).

Limitations. Three caveats. First, the BLP step is reported without asymptotic standard errors; bootstrap SEs will appear in the next revision. Second, the 12-point M4 over-prediction we interpret as incomplete transition rather than misspecification, but a full out-of-sample test requires post-2025 production data. Third, the counterfactual in Section 6 uses the firm-level $\widehat{\gamma}$ trajectory rather than re-solving

producer entry under no-cap; we report the welfare numbers as an accounting decomposition and defer the structural counterfactual.

Robustness of the structural estimates. The headline mechanism parameters are stable to two known modeling choices. Re-estimating on a grid of the regime-choice EV-scale $\sigma_{\text{reg}} \in [0.05, 1.00]$ leaves $\widehat{\varphi}$ saturated at 0.998 for $\sigma_{\text{reg}} \geq 0.15$ and $\widehat{\lambda}$ in $[0.58, 0.70]$/yr. Adding a multiplicative log-normal cost shock $\widetilde{C}_j = C_j \exp(\sigma_C \eta_j)$ with Monte Carlo integration over $N_{\text{sim}} = 50$ draws yields $\widehat{\sigma}_C = 0.038$ with every other estimate shifting by less than one percent versus the baseline. Detail in Appendix I.1.

6 Counterfactual and Robustness

6.1 No-Cap Counterfactual: How Much Loss Did the Cap Avert?

The bargaining-share identification in Section 5.2 prices the cap in firm-level RMB. Once $\widehat{\gamma}_{p,t}$ is in hand, a no-cap counterfactual is mechanical: hold γ_p at its pre-cap level from 2019 onward, keep all other firm-level variables (revenue, SG&A, R&D) at their realized values, and compute the implied operating income trajectory. Comparing this trajectory to the observed series yields a direct dollar measure of the cap's effect on the platform's bottom line.

Counterfactual specification. For each post-cap year $t \in \{2019, 2020, 2021, 2022\}$, the counterfactual operating income is

$$\widetilde{\Pi}_t \;=\; \text{Revenue}_t \cdot \gamma_p^{CF} \;-\; \text{SG\&A}_t \;-\; \text{R\&D}_t, \tag{17}$$

where γ_p^{CF} is the counterfactual bargaining share. The bracket below is anchored to two reference points from the model. The first is the Bertrand-auction outcome $\gamma_p^{\text{Bertrand}} = 0$ from Equation (12): in a no-cap world where the Bertrand auction mechanism is restored among K platforms, the star captures full marginal contribution and the platform's per-drama share is zero. The second is the 2018 trough $\widehat{\gamma}_{p,2018} = -0.086$, which adds the $(F + L)/R$ leverage burden on top of the Bertrand zero. Lifting the cap and assuming Bertrand conduct revives is a conduct/equilibrium-selection assumption, not a parameter-level prediction of the post-cap Nash model; the pre-cap industry behavior of Section 2.1 grounds this assumption in observed historical conduct. The three choices below span the resulting range:

- *CF-1:* $\gamma_p^{CF} = 0$. The content book exactly breaks even—the Bertrand-auction limit of Equation (12) under no cap, before adding leverage burden.
- *CF-2:* $\gamma_p^{CF} = -0.086$. The 2018 peak bidding-war intensity persists, adding the firm-level fixed and leverage burden $(F + L)/R \approx 0.09$ on top of the Bertrand zero; this is the maximally conservative no-cap scenario, cross-checked against the iQIYI 2018 trough.
- *CF-3:* $\gamma_p^{CF} = -0.035$. The 2016–2018 average $\widehat{\gamma}_p$, treated as the pre-cap baseline.

Result. Figure 4 plots the observed and three counterfactual operating-income paths. Table 13 reports year-level and cumulative gaps.

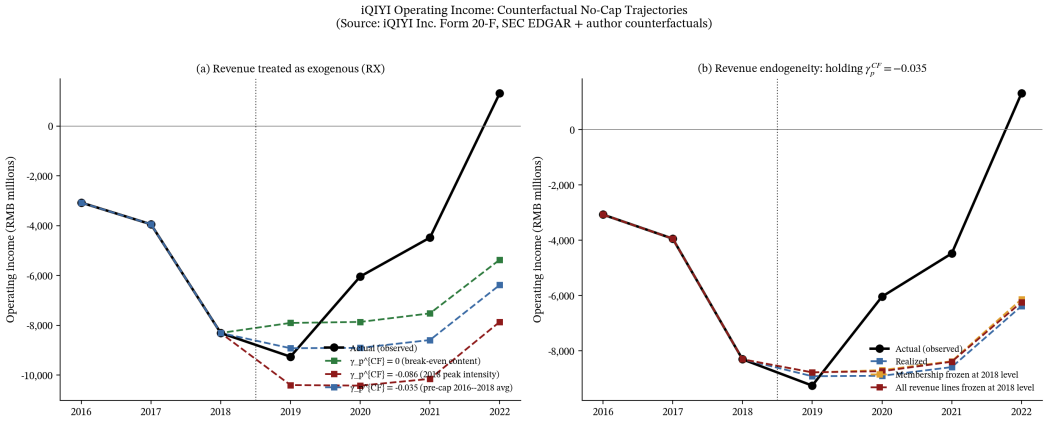


Figure 4: Observed (solid) vs. three counterfactual (dashed) iQIYI operating-income trajectories, 2016–2022, freezing γ_p at the indicated pre-cap value from 2019 onward with revenue and non-content overheads held at realized levels. Source: iQIYI 20-F.

Table 13: iQIYI No-Cap Counterfactual

Revenue counterfactual \ Bargaining share γ_p^{CF}	0 (break-even)	−0.086 (2018 peak)	−0.035 (pre-cap avg)
Realized	+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

Each cell reports the cumulative 2019–2022 gap $\sum_t(\Pi_t^{actual} - \widetilde{\Pi}_t)$ in RMB millions, under the indicated combination of revenue counterfactual (rows) and bargaining-share counterfactual (columns). $\widetilde{\Pi}_t = \text{Revenue}_t^{CF} \cdot \gamma_p^{CF} - \text{SG\&A}_t - \text{R\&D}_t$, with SG&A and R&D at realized values.
Realized: revenue treated as exogenous to the cap.
Membership frozen at 2018: subscription revenue stuck at its 2018 level (RMB 10.6B); ad and other revenues follow realized path. Reflects the view that subscription growth was cap-enabled.
All revenue frozen at 2018: total revenue stuck at RMB 25.0B. Maximally conservative revenue assumption for the no-cap world.
Positive cells indicate the cap raised cumulative operating income above the no-cap counterfactual. Across all $3 \times 3 = 9$ combinations the sign is positive, confirming that the cap-related recovery survives endogenous-revenue assumptions, though the magnitude varies by a factor of up to $\sim 4\times$.

The qualitative picture is unambiguous. Without the cap, even the most generous counterfactual (CF-1, zero content margin) leaves iQIYI losing RMB 5.4 billion in 2022 – 6.7 billion below the realized +1.3 billion. Under the 2018 peak-intensity continuation (CF-2), 2022 losses balloon to RMB 7.9 billion. The cumulative 2019–2022 gap between realized and counterfactual operating income ranges from RMB 10.2 billion (CF-1) to RMB 20.4 billion (CF-2). In all three specifications, the cap converts iQIYI from a chronically loss-making firm into one that crosses into operating profit in 2022.

Consumer-side accounting decomposition. Entry-margin dramas account for 60.1% of post-cap log-sum-exp consumer utility on average (49–65% year-by-year; Table 14). This is an accounting decomposition (lower bound on the structural welfare gain), not a re-solved counterfactual.

Table 14: Consumer Surplus: Entry-Margin Contribution

Year	Total dramas	Entry-margin dramas	CS observed (utils)	CS no-entry (utils)	Entry-margin share of CS
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: Per-market consumer surplus computed as the log-sum-exp over inside utility of the estimated BLP demand of Section 5.1. The observed CS uses all dramas with release year in $\{2019, \ldots, 2022\}$ in the estimation sample. The no-entry CS drops dramas directed by the locked-out director cohort identified in Section 3.1, holding all other primitives fixed. The entry-margin share of CS reports the fraction of observed consumer-side log-sum-exp utility attributable to entry-margin dramas, i.e., $1 - (\text{CS}_{\text{no-entry}}/\text{CS}_{\text{observed}})$.

This is a partial-equilibrium lower bound on the consumer-side welfare gain from the cap. It does not include the channel-expansion channel (Section C.4) or the quality-shift channel (Section C.2), each of which would further increase the consumer-side benefit. Because the predetermined-VA spec does not include price, the result is reported in utility units rather than RMB; a monetary translation requires the 2SLS price coefficient reported in the demand specification diagnostics.

Robustness and remaining caveats. Revenue is partly endogenous to the cap. A 3×3 grid that adds two revenue counterfactuals (RY freezes subscription revenue at 2018; RZ freezes total revenue at 2018) yields a positive cumulative saved-loss estimate in every cell, with the conservative 10–15 billion range (RY/RZ revenue, $\gamma_p^{CF} \in \{0, -0.035\}$) as the defensible point estimate. The partial-equilibrium frame still abstracts from cross-platform spillovers.

6.2 Structural No-Cap Counterfactual: Re-Solving Each Player

The accounting counterfactual of Section 6.1 holds the bargaining share at a pre-cap reference value and reads the implied dollar gap off iQIYI's 20-F line items. Because the cap is identified, in this paper, as a switch in the platform–talent regime (Bertrand to Nash) and a restart of the capability accumulation path, the structural model permits a sharper exercise: re-solve every agent's choice under the no-cap regime and compare the implied equilibrium objects to their observed counterparts. Three structural objects are re-computed—platform make-or-buy share, director entry payoff, actor career-path value—using the estimates of Sections 5.1–5.5.

No-cap regime. The structural CF fixes the two primitives that the cap moved:

$$\varphi^{\mathrm{CF}} = 1, \qquad \Delta\kappa^{\mathrm{CF}}(t) = \Delta\kappa_0 \quad \text{for all } t \geq 2018, \tag{18}$$

i.e., the Bertrand auction remains binding ($d_S^{\mathrm{CF}} = V_a$, $\gamma_M^{\mathrm{CF}} = 0$) and the make–buy wedge does not decay (no capability buildup post-2018). All demand parameters (α, β, Σ), the Nash weight ω, the entry primitives κ_d, and the actor flow-utility parameters (f_Q, δ_Q, β) are held at their estimated values. Every post-2018 equilibrium object is then re-derived from these primitives exactly as in Section 5.3.

Result A: make-or-buy share collapses. Under $\varphi^{\mathrm{CF}} = 1$, the platform's make-side surplus $\gamma_M^{\mathrm{CF}}(j,t) = \omega(1 - V_a/V_a)_+ = 0$ for every project: the Bertrand auction transfers the entire star surplus to talent. Holding the buy-regime cost wedge at $\Delta\kappa_0 = 2.054$ also makes make weakly dominated in cost. The model-implied counterfactual make share stays in the ∼2–10% band throughout 2019–2024, against an observed series rising from 20% to 48% (Figure 5, panel (a)). The vertical-integration transition documented as Fact 4 in Section 3 is entirely off-equilibrium without the cap.

Result B: director entry payoff stays low. The per-project expected payoff $E[\max(\pi_M, \pi_B)]$ that drives the director's binary-choice problem (Equation (5)) is held down by the same two forces. Under the no-cap CF, $E[\max(\pi_M, \pi_B)]$ stays below 0.05 throughout 2019–2024, against an observed rise from 0.175 in 2019 to 0.376 in 2023. Plugging these payoff trajectories into the estimated probit yields a counterfactual locked-out cohort entry rate of ∼13% throughout post-2018 (essentially the pre-cap rate), against the observed 44.6% for streaming incumbents and the −22.6 pp average marginal effect of Section 5.4. The director-entry margin documented as Fact 3 does not appear in the no-cap world.

Result C: actor welfare and the transfer accounting. Re-solving the actor's Bellman (Equation (8)) under the regime-weighted wage schedule that prevails in the CF (still close to $W = V_a$ for top stars) yields a value function V_a^{CF}. The observed-minus-CF gap $V_a^{\mathrm{obs}} - V_a^{\mathrm{CF}}$ is negative (−6.0 utility units, averaged across states): the cap reduces incumbent top-star lifetime value by transferring ∼6 units of welfare from actors to producers and consumers. The accounting per-drama profit-gap series $\pi_{j,t}^{\mathrm{obs}} - \pi_{j,t}^{\mathrm{CF}}$ is positive throughout 2017–2021 and turns negative only after 2022, with a cumulative post-2018 net of +0.33 per drama in favor of the observed regime—consistent with the platform-side surplus recovery documented in Fact 1.

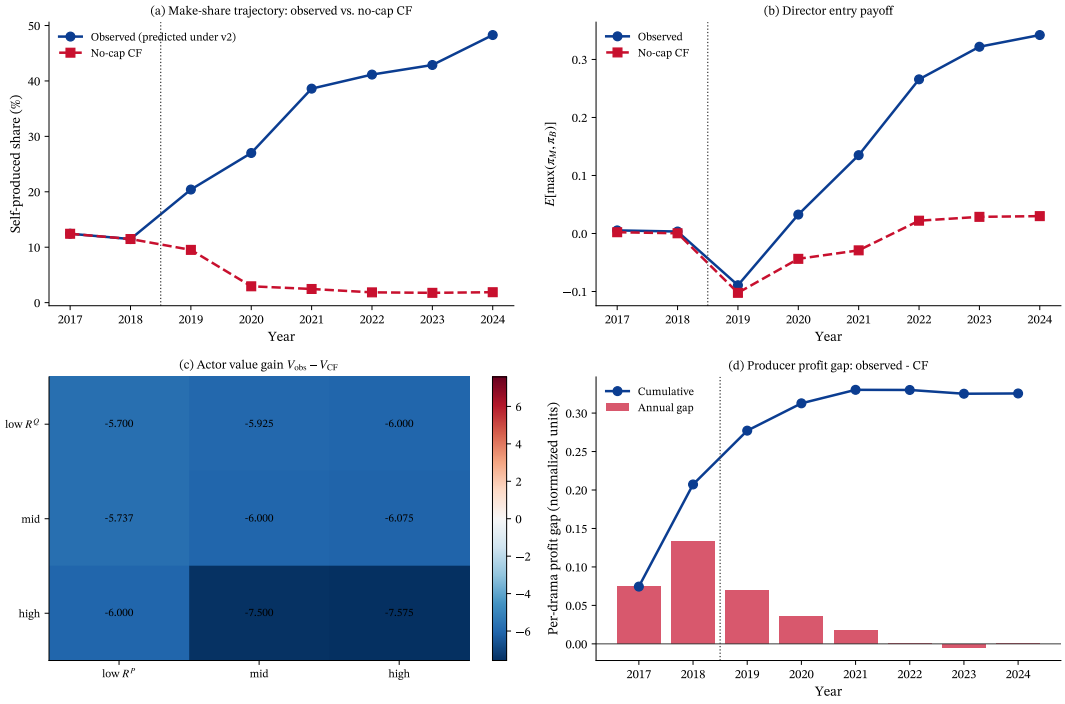


Figure 5: Structural no-cap counterfactual under $\varphi^{\text{CF}} = 1$ and $\Delta\kappa^{\text{CF}}(t) = \Delta\kappa_0$. (a) Make-or-buy share: observed ($20 \to 48\%$) vs. CF (flat $\sim$2–10%). (b) Director entry payoff $E[\max(\pi_M, \pi_B)]$: observed ($0.175 \to 0.376$) vs. CF (flat ≤ 0.05). (c) Actor Bellman gain $V^{\text{obs}} - V^{\text{CF}}$: uniformly ≈ -6 across states $\Rightarrow$ cap transfers welfare from actors. (d) Cumulative per-drama profit gap: +0.33 in favor of the observed regime.

Accounting vs. structural CF. The accounting CF (§6.1) asks what dollar profit iQIYI would have realized had γ_p stayed at the pre-cap reference (RMB 10–20 B saved); the structural CF asks whether the post-2018 entry, variety, and integration margins exist at all without the regime switch (no—under $\varphi^{\text{CF}} = 1$ and $\Delta\kappa^{\text{CF}} = \Delta\kappa_0$ the platform stays in the pre-cap corner solution). The entry-margin dramas contribute 60% of post-cap consumer-side utility (Table 14).

Ruling out alternative interpretations. The headline transition is not a pre-existing trend, a COVID demand shock, a generic global subscription-streaming pivot, or a pure wage compression. Pre-cap linear-trend extrapolation of $\widehat{\gamma}_{p,t}$ predicts -20.5% in 2022 vs. the realized $+23.0\%$ (a 43.5 pp structural break); the entry,

variety, and vertical-integration shifts begin in 2019, pre-pandemic; Netflix exhibits no analogous break over the same window; and the make–buy, entry, and matching margins move together as the regime-switching model predicts and the pure-wage-compression alternative does not. Detail in Appendix J.1.

7 Conclusion

The 2018 Chinese actor salary cap functioned not as a marginal labor-cost intervention but as the lever that broke single-content competition among long-form streaming platforms. By rendering a binding ceiling on above-line compensation enforceable, the policy severed the Bertrand-auction-driven loop between traffic stars, advertising revenue, and the production-cost floor that sustained the pre-cap monoculture. Structurally, the pre-cap auction was at near-saturation ($\widehat{\varphi} = 0.998$); the cap collapsed it and launched capability accumulation at $\widehat{\lambda} = 0.64$/yr, yielding a $\sim$70% structural wage compression and admitting the narrative-genre directors who had been incentive-incompatible with the prior regime. Entry-margin dramas account for 60% of post-cap consumer-side log-sum-exp utility under the estimated BLP demand; the accounting counterfactual implies RMB 10–20 billion in averted iQIYI losses over 2019–2022.

Generalization. The mechanism is not specific to China. Wherever an upstream input has tournament-shaped rewards and downstream channels concentrate on it, the market can lock into a monocultural production technology that excludes whole classes of producers; a regulatory intervention on the input price unlocks the excluded technologies, with the welfare gain operating through consumer variety rather than rent redistribution. The same primitive structure appears in U.S. long-form streaming (the 2023 SAG-AFTRA settlement), the mid-2000s music industry (digital distribution collapsing the record-label A&R tournament, Aguiar and Waldfogel, 2018, Waldfogel, 2018), and plausibly in news media and academic publishing where no enforceable instrument has unlocked the channel. The Chinese cap is one realization of a regulatory category—public input-price intervention as a channel-monoculture remedy—that, to our knowledge, has not been previously identified in the empirical industrial-organization literature.

References

Moshe Adler. Stardom and talent. American Economic Review, 75(1):208–212, 1985.

Luis Aguiar and Joel Waldfogel. Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music. Journal of Political Economy, 126(2):492–524, 2018.

Simon P Anderson and Stephen Coate. Market provision of broadcasting: A welfare analysis. The review of Economic studies, 72(4):947–972, 2005.

Peter Arcidiacono and Robert A Miller. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. Econometrica, 79(6):1823–1867, 2011.

Mark Armstrong. Competition in two-sided markets. RAND Journal of Economics, 37(3):668–691, 2006.

Susan Athey and Kyle Bagwell. Collusion with persistent cost shocks. Econometrica, 76(3):493–540, 2008.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. Econometrica, 63(4):841–890, 1995.

Steven Berry, Martin Gaynor, and Fiona Scott Morton. Do increasing markups matter? lessons from empirical industrial organization. Journal of Economic Perspectives, 33(3):44–68, 2019.

Steven T. Berry and Joel Waldfogel. Free entry and social inefficiency in radio broadcasting. Quarterly Journal of Economics, 114(2):397–420, 1999.

Russell Cooper, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. Selection criteria in coordination games: Some experimental results. American Economic Review, 80(1):218–233, 1990.

Gregory S. Crawford, Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu. The welfare effects of vertical integration in multichannel television markets. Econometrica, 86(3):891–954, 2018.

Helmut M Dietl, Egon Franck, Markus Lang, and Alexander Rathke. Salary cap regulation in professional team sports. Contemporary Economic Policy, 30(3): 337–349, 2012.

Robert H Frank and Philip J Cook. 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. 1995.

Bruce C. Greenwald and Joseph E. Stiglitz. Externalities in economies with imperfect information and incomplete markets. Quarterly Journal of Economics, 101(2):229–264, 1986.

Kate Ho and Robin S. Lee. Insurer competition in health care markets. Econometrica, 85(2):379–417, 2017.

Bengt Holmström. Moral hazard and observability. Bell Journal of Economics, 10 (1):74–91, 1979.

V Joseph Hotz and Robert A Miller. Conditional choice probabilities and the estimation of dynamic models. The Review of Economic Studies, 60(3):497–529, 1993.

Lawrence M. Kahn. Cartel behavior and amateurism in college sports. Journal of Economic Perspectives, 21(1):209–226, 2007.

N. Gregory Mankiw and Michael D. Whinston. Free entry and social inefficiency. RAND Journal of Economics, 17(1):48–58, 1986.

Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. Econometrica, 69(2):307–342, 2001.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. Journal of the European Economic Association, 1(4):990–1029, 2003.

Sherwin Rosen. The economics of superstars. The American Economic Review, 71 (5):845–858, 1981.

A. Michael Spence. Product selection, fixed costs, and monopolistic competition. Review of Economic Studies, 43(2):217–235, 1976.

John Sutton. Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration. MIT Press, Cambridge, MA, 1991.

Joel Waldfogel. Preference externalities: An empirical study of who benefits whom in differentiated-product markets. RAND Journal of Economics, 34(3):557–568, 2003.

Joel Waldfogel. Digital Renaissance: What Data and Economics Tell Us about the Future of Popular Culture. Princeton University Press, 2018.

A Data Appendix: Platform Revenue and Production Strategy

A.1 Verification of Platform Revenue Data

iQIYI revenue data come directly from audited Form 20-F filings with the SEC. Tencent Video and Youku are reported only at the parent-company segment level (Tencent Holdings, Alibaba) but management discussions corroborate the same industry-wide trend.

iQIYI revenue composition inverts sharply around 2018: advertising falls from 46.9% (2017) to 18.3% (2022), while membership rises from 37.6% to 61.0% (Form 20-F filings, 2018/2020/2022). Operating losses peak at RMB 9.3 billion in 2019 and narrow steadily through 2022, the firm-level realization of the cost-side recovery captured by $\widehat{\gamma}_{p,t}$ in Figure 2.

A.2 The Economic Rationale for Pre-Regulation Outsourcing

Pre-cap, with revenue tied to ad-driven viewership and traffic delivered by top stars, licensing a star-vehicle from an external studio was the dominant strategy: it transferred production risk, delivered a pre-packaged marketing asset, and let platforms focus on distribution and ad sales. The cap removed the traffic-star input, broke the advertising model, and catalyzed the pivot to subscription. Under subscription, the objective shifts to long-term subscriber value, which depends on exclusive high-quality content—making vertical integration the new dominant strategy (quality control, exclusive IP, restored bargaining power). This is the structural counterpart to the induced-vertical-integration mechanism of Hart-Tirole (1990).

B Construction of Star Value Added

We construct actor-level value-added measures from each actor's prior works, using both Douban rating (rating-VA) and log 30-day viewership (views-VA) as the underlying quality signal. Let j index works, a index actors, t index calendar years. For $X \in \{\text{rating}, \text{views}\}$ and $Q_j^X \in \{\text{Douban score}_j, \log v_{30,j}\}$, year-demeaned quality is

$$\widetilde{Q}_j^X = Q_j^X - \bar{Q}_{y(j)}^X,$$

removing common year shocks. Each actor–work observation is weighted by billing order b_{aj}:

$$w_{aj} = \begin{cases} 1.0, & b_{aj} \in \{1, 2\}, \\ 0.8, & b_{aj} \in \{3, 4, 5\}, \\ 0.5, & b_{aj} \in \{6, 7\}, \\ 0.3, & b_{aj} \in \{8, 9, 10\}, \\ 0.1, & \text{otherwise or missing.} \end{cases}$$

Actor a's value-added in year t is the billing-weighted average of pre-t year-demeaned quality:

$$\mathrm{VA}^X_{at} = \frac{\sum_{j:a\in j,\, y(j)<t} w_{aj}\, \widetilde{Q}^X_j}{\sum_{j:a\in j,\, y(j)<t} w_{aj}},$$

with missing Q^X_j excluded from both sums. The pre-2018 snapshot $\mathrm{VA}^X_{a,2018}$ is used as the predetermined cast-side regressor in BLP demand (Section 5.1) and as the state variable in the actor's CCP estimation (Section 5.5).

C Detailed Reduced-Form Evidence

C.1 Cross-Medium Director Reallocation

C.2 Equilibrium Transition: Who Re-Entered the Market

The cap's visible margin of adjustment lies on the supply side: audience preferences do not noticeably shift, but who participates in streaming-drama production does. Two groups respond—(i) high-reputation incumbent actors substituting toward endorsement income, and (ii) experienced directors previously locked out of streaming by the dominance of star-vehicle projects. Tables 15 and 16 quantify both.

Who is shifting away from acting? Table 15 restricts to the 3,298 actors active in both 2017–2018 and 2019–2022 and decomposes their own career shift by pre-2018 rating-reputation tertile. Three findings appear together. First, low- and mid-reputation actors shift weakly toward (or are flat on) high-quality lead roles (+0.87pp and +0.09pp on ΔP(lead high)), while high-reputation actors lose 2.75pp of their lead-high share and gain 2.07pp of lead-low share. Second, brand endorsements rise across all reputation bins, but with the opposite ordering: low-bin actors gain +0.80 endorsements per actor on average; high-bin actors gain only +0.26. Third, and most diagnostic, the within-actor correlation between endorsement growth and

high-quality lead-role acquisition flips sign across bins. For low- and mid-bin actors, $\rho(\Delta\text{endor}, \Delta\text{lead high}) = +0.161$ and $+0.240$ respectively—endorsements and high-quality leads rise together, consistent with rising stars accumulating both. For high-bin actors, $\rho = -0.182$ and $\rho(\Delta\text{endor}, \Delta\text{lead low}) = +0.351$: those high-reputation actors who pick up more endorsements take fewer high-quality leads and more low-quality leads. The substitution at the high-reputation margin is direct: brand contracts displace prestigious acting work, not the reverse.

Who is supplying post-policy quality content? Table 16 addresses the corresponding director-side question. Of the 487 directors who helmed a streaming drama between 2019 and 2022, 401 were absent from the 2017–2018 sample; we call these *post-only* directors. Panel A demonstrates that they are not new entrants. Only 32 of 401 (about 8%) have no pre-2017 Douban filmography; another 34 debuted in the immediate pre-policy years 2015–2016. The bulk are *Mid-career* directors who started 2010–2014 ($n = 60$), *Veterans* from 2000–2009 ($n = 125$), and *Legacy* directors whose first credit predates the year 2000 ($n = 150$). The median post-only director has 11 prior works and a debut around 2002.

Timing rules out a "return from retirement" interpretation. Among the 369 post-only directors with any pre-2017 Douban-recorded work, the median latest pre-2017 work year is 2016 (25th percentile 2015), an average gap of ≈ 1.9 years between last pre-cap credit and first post-cap streaming credit. These directors remained active in adjacent media—broadcast television, theatrical film, lower-budget online production—and reallocated into streaming after the cap relaxed the production-cost floor.

Panel B compares the pre-2017 Douban score of post-only directors to the 86 "both-period" streaming incumbents: 2.37 vs. 2.49 on the five-point proxy, a -0.12 gap with 95% Welch CI $[-0.32, +0.08]$ ($t = -1.22$, $p = 0.22$); the comparison is underpowered ($n = 401$ vs. 86) so a modest gap in either direction is not ruled out. The economic claim that post-only directors are experienced crowded-out producers rather than novice entrants rests on the Panel A career-stage evidence.

Panel C closes the loop. Among post-2018 high-quality dramas (Douban ≥ 7), **81.9**% are directed by post-only entrants, against 18.1% by both-period incumbents. The directors supplying the post-policy quality of the medium are overwhelmingly those previously locked out of it.

Both adjustments—incumbent actors reallocating effort into endorsements, and experienced directors entering streaming from adjacent media—realize the supply-side equilibrium transition formalized in Section 4 and Appendix E.

Table 15: Within-Actor Career Shift

		Drama composition (pp)		Endorsements			
Pre-2018 rating bin	N	ΔP(lead high)	ΔP(lead low)	Pre (avg)	Post (avg)	Δ	$\rho(\Delta\text{endor}, \Delta\text{lead_high})$
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 lead-high roles by 1.6–3.4pp and gain lead-low roles by 1.9–2.4pp, while their endorsement-substitution correlation flips sign ($\rho < 0$) relative to lower-reputation actors ($\rho > 0$): high-reputation actors who land more endorsements take fewer high-quality lead roles.

Table 16: Origin of Post-Policy Director Entrants

Panel A: Career stage of post-only directors		
Career stage at first credit	N	Share
Truly new	32	8.0%
Recent (2015-16 debut)	34	8.5%
Mid-career (2010-14)	60	15.0%
Veteran (2000-09)	125	31.2%
Legacy (pre-2000)	150	37.4%
Panel B: Pre-2017 reputation, post-only vs. both-period		
Average pre-2017 Douban score (post-only)		2.374
Average pre-2017 Douban score (both-period)		2.494
Difference (post-only − both)		-0.121
95% CI on difference (Welch)		$[-0.316, +0.075]$
Welch two-sided t-test p-value		0.225
TOST equivalence margin $\delta = \pm 0.25$ Douban points		
TOST p-value (one-sided max)		0.097
Equivalence established at 5%?		No
Panel C: Who supplies post-2018 high-quality dramas (Douban ≥ 7)?		
Total post-2018 high-quality dramas		155
Share led by both-period directors		18.1%
Share led by post-only directors		81.9%

Post-only directors are those credited on a 2019–2022 drama but absent from 2017–2018. *Both-period* directors are credited in both windows. Career stage is the year of the director's first credit in the Douban filmography panel (`20_director_creations`).

Panel B uses the Douban rating-class proxy (`allstarNN`/10). The Welch t-test does not reject equality of pre-2017 reputation between post-only and both-period directors. To make a positive statement about equivalence, we report a TOST procedure with margin $\delta = 0.25$ Douban points (5% of the 5-point scale). When the TOST p-value is below 0.05, the difference is statistically equivalent to zero within $\pm\delta$; otherwise we cannot conclude equivalence and report only that we fail to detect a quality gap (the test is underpowered).

Panel C demonstrates that the bulk of post-policy high-quality streaming dramas is supplied by directors who were not in the streaming market in 2017–2018, despite holding comparable pre-policy reputations.

C.3 Revenue-Channel Expansion and Market Entry

C.4 Revenue-Channel Expansion and Market Entry

The post-cap differentiation strategy predicts coexistence of the pre-cap advertising-funded star-vehicle tier with a new, subscription-funded long-tail tier rather than replacement. This subsection documents the coexistence with two auxiliary data sources.

Channel expansion at the drama level. We use drama-level brand-sponsorship records from Endata covering 715 dramas with on-air dates 2017–2023. For each drama we observe the number of sponsoring brands (品牌数) and total brand on-screen exposure (品牌露出总时长 (秒)). Aggregating to the year level, Table 17 reports the distribution and Figure 6 the two key trajectories.

The data do not support a uniform decline in advertising activity. The maximum brand count per drama is volatile across years (122 in 2018, 109 in 2020, 51 in 2022): genuine post-cap blockbusters such as Thirty's Up (《三十而已》) continue to attract concentrated sponsorship at the high end. The advertising-driven business model remains operational. What changes is the lower tail of the distribution: the share of dramas with ≤ 5 sponsoring brands roughly doubles, from 13–16% in 2017–2018 to 32–34% in 2020–2022. The top-10 share of total brand-exposure also falls (from $\approx 70\%$ in 2017–2018 to 35–50% post-cap), but its trajectory is non-monotonic.

The pattern is most naturally read as the emergence of a parallel, subscription-funded content tier. These low-brand-density dramas are not the same projects that would previously have competed in the advertising tournament; instead, they are new productions whose cost structure (no longer dominated by top-star compensation) makes them viable to recover costs through paid subscriptions and to operate alongside, rather than against, the existing high-budget ad-funded tier. This interpretation aligns with the realized λ_t trajectory in Section 4's platform profit function, and with iQIYI's documented revenue-mix inversion (membership 37.6% $\rightarrow$ 61.0%; advertising 46.9% $\rightarrow$ 18.3%, 2017–2022, Form 20-F).

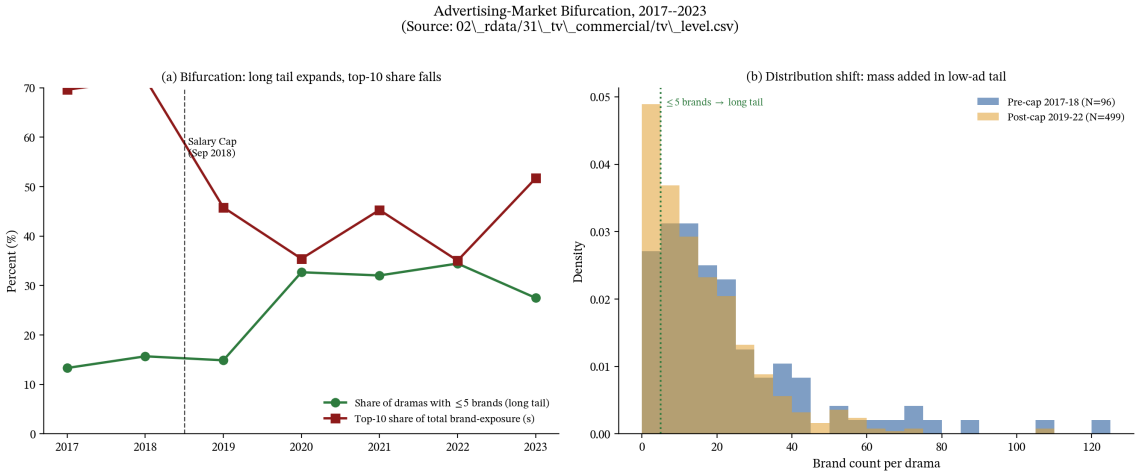


Figure 6: Advertising-Market Bifurcation

Notes: Panel (a) plots the long-tail share (dramas with ≤ 5 sponsoring brands) and the top-10 share of total brand-exposure time. Panel (b) compares the density of brand-count per drama in the pre-cap (2017–2018) and post-cap (2019–2022) periods. Source: Endata drama-level brand-sponsorship records, $N = 715$ dramas. *Reading:* Post-cap mass is added in the low-end tail while the high-end tier remains populated – the empirical signature of the new low-cost channel emerging alongside (not displacing) the advertising-funded tier, as predicted by the two-front squeeze differentiation strategy of Section 1.

Table 17: Advertising-Market Bifurcation

		Distribution of brand count (品牌数)		Tails		Concentration	
Year	N	Median	90%ile	Long tail[a]	Top tier[b]	Gini	Top-10 share[c]
2017	45	20.0	41.6	13.3%	6.7%	0.380	69.6%
2018	51	15.0	66.0	15.7%	15.7%	0.530	71.8%
2019	121	17.0	38.0	14.9%	5.8%	0.410	45.8%
2020	156	11.5	38.0	32.7%	7.7%	0.530	35.4%
2021	103	13.0	34.0	32.0%	1.9%	0.460	45.3%
2022	119	10.0	24.0	34.5%	0.8%	0.440	35.0%
2023	120	12.0	36.0	27.5%	0.8%	0.440	51.7%

Source: the brand-sponsorship panel (Endata), drama-level data covering 1,133 dramas with on-air dates 2016–2026. Brand count (品牌数) is the number of distinct sponsoring brands appearing on the drama; total brand exposure is the cumulative on-screen exposure time across all sponsors.

[a] Share of year's dramas with ≤ 5 sponsoring brands (the empirical long tail).

[b] Share of year's dramas with ≥ 50 sponsoring brands (top tier).

[c] Concentration of total brand-exposure on the top-10 dramas of each year.

Key readings: (i) the long-tail share doubles from $\approx 15\%$ in 2017–18 to 32–34% in 2020–22; (ii) the top-10 share of brand exposure collapses from 58% in 2018 to $\approx 25\%$ post-cap; (iii) the maximum brand count is not monotone (2018 peak 122, 2020 still 109)—genuine blockbusters such as Thirty's Up continue to capture concentrated sponsorship. The pattern is one of **market bifurcation**—a long tail of low-sponsorship productions emerges below a still-present top tier—rather than a uniform collapse of the advertising tournament.

Product-variety expansion on the supply side. The channel-coexistence finding above is mirrored in the cast-and-crew composition of the post-cap drama supply. Holding annual broadcast volume approximately constant (around 430 dramas per year), Figure 7 traces four variety measures at the release-year level. The count of distinct directors active in any given year more than doubles from 97 in 2018 to a 2019–2022 average of 192, so the average drama-per-director falls from 4.8 to 2.2. The Herfindahl on director shares of within-year drama production halves (from 0.012 to 0.006), and the analogous concentration on top-billed actors falls in parallel. The expansion is a direct, descriptive counterpart of the entry-margin DiD of Section 3.1: more producers reach the streaming-drama market each year, on tighter average loads, post-cap. Combined with the BLP demand estimates of Section 5.1, this supply-side variety expansion produces the consumer-side log-sum-exp welfare gain documented in Section 6.

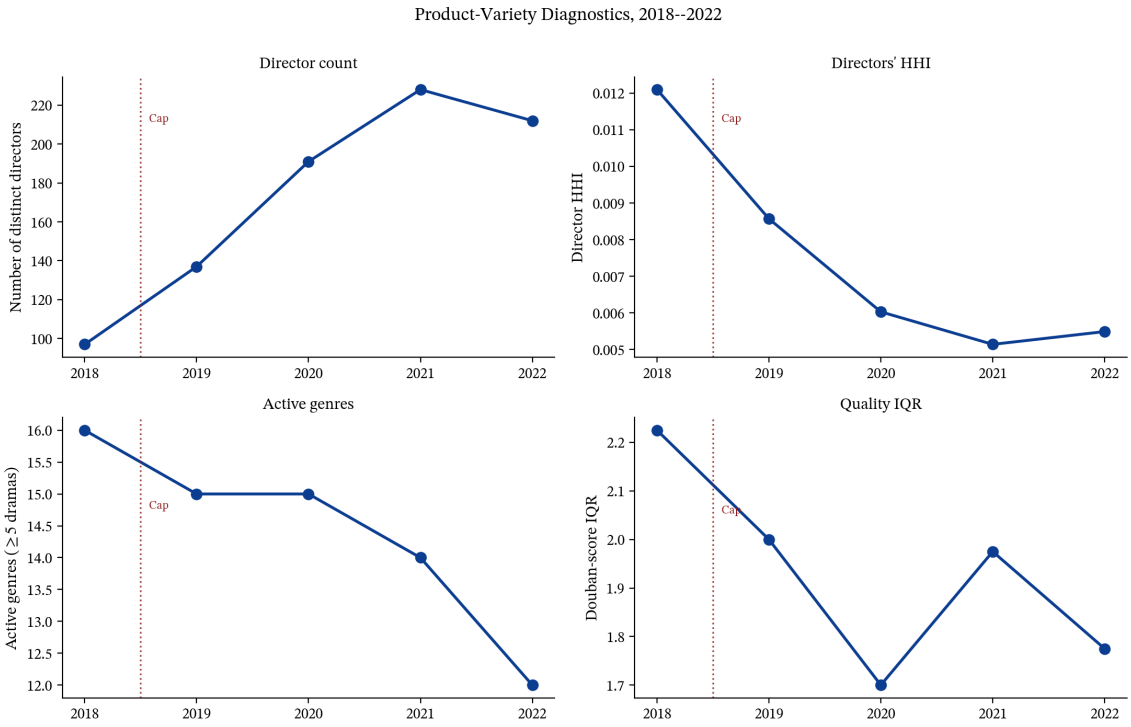


Figure 7: Product-Variety Diagnostics

Notes: Year-level variety diagnostics on the $N \approx 2{,}170$ drama panel for 2018–2022. Top-left: count of distinct directors. Top-right: director Herfindahl ($\sum_i q_i^2$ where q_i is director i's share of within-year dramas). Bottom-left: number of Douban genres with at least five dramas. Bottom-right: interquartile range of Douban scores within the year. Source: drama panel (Yunhe-Douban-Endata). *Reading:* Top-row panels show the central variety expansion: the count of active directors more than doubles and director HHI halves, both inside three years post-cap. Bottom-row panels are ancillary; the active-genre count is approximately flat and Douban IQR narrows somewhat, consistent with a quality-distribution compaction rather than further dispersion.

Genre composition and the idol-drama filter. The variety expansion is not uniform across content types: post-cap supply shifts decisively from spectacle and star-vehicle genres to narrative-driven ones. The largest gainers in within-year share are Crime (+3.9pp, nearly doubling from 4.1% to 8.0%), Drama (+3.1), History (+1.3, +63% relative), Mystery (+1.2), and Costume Drama (+0.8). The largest losers are Comedy (−6.7pp, −35% relative), Fantasy (−3.3), Sci-Fi (−2.0, −60%), Action (−1.7), and Adventure (−1.0). Decomposing each shift into an entry-composition channel (entry directors having a different genre type than incumbents) and an

incumbent-transition channel reveals that 88–94% of the gain in the rising narrative genres comes from entry composition: entering directors specialize in different content types rather than incumbents reinventing themselves. Detailed tables and the full genre ranking are available from the authors.

These patterns map directly onto the idol-drama (偶像剧) filter: pre-cap, only formats compatible with absorbing the top-star fraction in a single budget line item got made (romance, fantasy, comedy, costume drama), and narrative-driven directors (the 正剧 tradition—crime, mystery, history) were excluded by the cost floor. The cap relaxes the bottleneck and admits this previously-excluded segment. The 60 percent of post-cap consumer log-sum-exp utility attributable to entry-margin dramas (Section 6) is therefore approximately the welfare contribution of a genre–director combination that the pre-cap business model structurally excluded.

Market entry on the supply side. We complement the broadcast-level evidence with an announcement-level view from iQIYI's publicly-disclosed drama pipeline (the Tencent and Youku rows of the same source lack the announcement-year coverage required for a time-series read).[3] Of 400 iQIYI announcements with parsable year, 32 are pre-cap (2015–2018) and 367 are post-cap (2019–2025). Two cross-checks.

1. *Production-volume growth.* iQIYI announcements scale from 1–18 per year pre-cap to 49–62 per year post-cap (Figure 8, panel a), consistent with the broadcast-volume expansion reported in Section 3.

2. *Lead-pool widening.* Each announcement lists approximately four named leads. The 32 pre-cap announcements draw on a pool of 111 unique individuals; the 367 post-cap announcements draw on a pool of 765 unique individuals, of whom **687**—fully 90%—are entirely new relative to the pre-cap pool. This is the actor-side analog of the director-side market-entry finding in Table 16: the post-cap streaming market draws on a far wider, far broader pool of named talent than the pre-cap market did, suggesting that the cap relaxed a binding budget constraint that had previously excluded the bulk of working actors from leading roles in this medium.

[3]This is a sample of platform-announced projects, not the universe of iQIYI productions; the levels are not directly comparable to the broadcast panel.

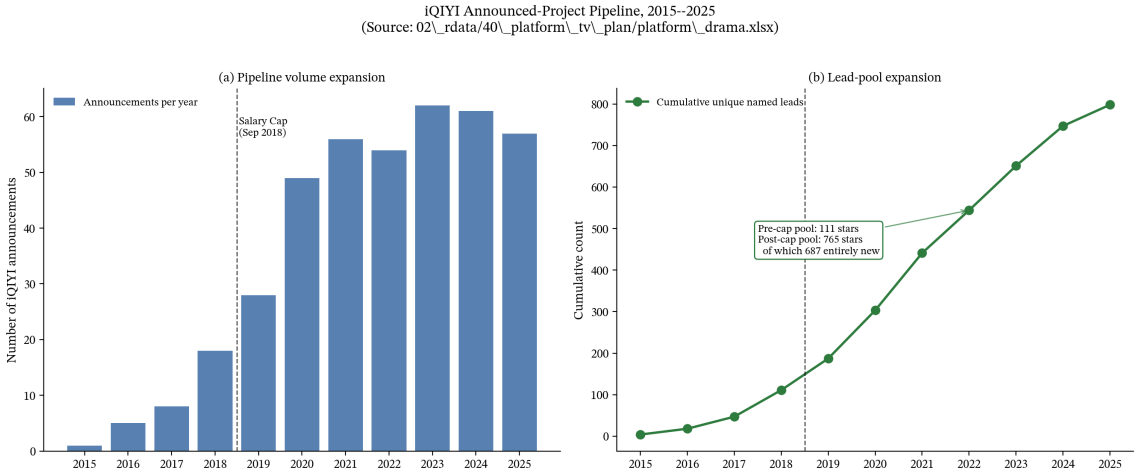


Figure 8: iQIYI Project Pipeline and Lead-Pool Expansion

Notes: Panel (a) plots the number of annual project announcements. Panel (b) plots the cumulative count of unique named leads across all announcements through the indicated year. Source: iQIYI publicly-disclosed announced-project pipeline, 2015–2025. *Reading:* The post-cap cumulative lead pool is 6.9× larger than the pre-cap pool, with 687 of the 765 post-cap leads not appearing in any pre-cap announcement. This is the actor-side realization of the entry margin documented for directors in Section 3.1: a broader pool of working talent gains streaming access after the cap relaxes the production-cost floor.

Table 18: iQIYI Project Pipeline and Lead-Pool Expansion

Panel A: annual pipeline

Year	Announcements	Lead slots	Unique leads	Leads/drama	New leads vs. cumulative
2015	1	4	4	4.00	4
2016	5	20	16	4.00	14
2017	8	32	32	4.00	29
2018	18	72	68	4.00	64
2019	28	112	100	4.00	76
2020	49	195	172	3.98	116
2021	56	224	199	4.00	138
2022	54	216	198	4.00	103
2023	62	246	223	3.97	107
2024	61	243	217	3.98	96
2025	57	223	187	3.91	51

Panel B: pre vs. post lead pool

Pre-cap announcements (2015–2018)	32 dramas	111 unique leads
Post-cap announcements (2019–2025)	367 dramas	765 unique leads
of which in both periods (overlap)		78
of which **new** post-cap		**687**

Source: iQIYI publicly announced drama production pipeline, by announcement year. (Tencent $N = 183$ and Youku $N = 103$ entries lack reliable announcement-year coverage and are excluded.)

Column definitions (Panel A). **Announcements**: count of distinct drama projects iQIYI announced as in production in that year. **Lead slots**: total number of billed lead positions across the year's announcements, summed over dramas. Each drama lists ≈ 4 billed leads by industry convention. **Unique leads**: count of distinct named actors appearing in the 主演 (lead-actor) field of the year's announcements; an actor billed in two dramas the same year counts once. **Leads/drama**: lead slots $\div$ announcements (≈ 4 throughout, confirming the billing convention is stable). **New leads vs. cumulative**: number of leads in the year who did not appear in any earlier year's iQIYI announcement (i.e., the gross flow of new lead-actor entrants into iQIYI's pipeline that year).

Panel B definitions. **Pre-cap pool**: the set of distinct actors appearing as a lead in any iQIYI announcement during 2015–2018. **Post-cap pool**: same definition over 2019–2025. **Overlap**: actors appearing in both pre-cap and post-cap pools (i.e., pre-cap leads who continued to be cast as leads post-cap). **New post-cap**: post-cap leads who never appeared as a lead in any iQIYI announcement during 2015–2018. By construction, New post-cap = Post pool − Overlap.

Reading: post-cap iQIYI announces 11.5× more dramas (32 vs. 367) but draws on a 6.9× larger pool of distinct lead actors (111 vs. 765). Of the 765 distinct lead actors in the post-cap pool, **687** (90 percent) had never appeared as a lead in any iQIYI announcement before September 2018—they are entirely new lead-actor entrants to the platform. This is the actor-side counterpart to the director-side entry margin documented in Table 16.

Implication: differentiation, not channel substitution. The pattern is differentiation rather than substitution: the advertising channel remains operational for high-budget productions of the pre-cap variety, while a subscription channel becomes independently viable for low- to mid-budget productions previously infeasible under the cost floor. The joint effect is positive-sum across viewers (content volume 4–10×, lead pool 6.9× wider, no withdrawal of existing options), entering producers and crowded-out talent (687 new leads plus entry-margin directors), and platforms (the $\widehat{\gamma}_p$ recovery from -8.6% to $+23\%$ is the margin on the new channel, not rent recovered from incumbent top talent whose post-cap project participation is largely unchanged in volume).

Caveats. Both data sources are partial. The brand-sponsorship panel (Endata) is a sponsorship registry whose coverage of small productions may be selected (a $\leq$ 5-brand drama could be a true long-tail production or an under-recorded entry; the level of the long-tail share is therefore noisy though the doubling is robust). The platform-pipeline source is an announcement registry, not a broadcast panel.

D Identification of the Structural Primitives

This appendix states the identifying moment for each deep parameter recovered in Section 5's pipeline.

D.1 Demand: $(\alpha, \beta, \Sigma, \xi_j)$

Standard BLP, identifying restrictions detailed in Section 5.1. Predetermined-VA spec uses the actor's pre-2018 work history $R^P_{a,<2018}$ as the cast-side regressor, removing the post-policy casting-selection channel from identification. BLP-style instruments (competitor-characteristic sums) discipline the price coefficient α.

D.2 Bargaining Weight ω

The bargaining weight ω is identified from the post-cap non-binding subsample. We use the disclosed $N = 10$ iQIYI projects whose actor-cost share lies strictly below the 40% C1 ceiling—i.e., where the cap constraint is slack and the Nash FOC of Equation (14) holds with equality:

$$W_a \;=\; (1-\omega)V_a \;+\; \omega d_S.$$

Rearranging,

$$\omega \;=\; \frac{V_a - W_a}{V_a - d_S}. \tag{19}$$

Source of variation. Equation (19) contains three observables and one structural parameter:

- V_a: BLP-derived marginal revenue (from §D.3 below).
- W_a: observed lead-actor wage, from the $N = 10$ 主演成本 disclosure rebased to revenue via the iQIYI year-specific C/R ratio.
- d_S: calibrated outside option (§D.4).

ω is point-identified from any single (V_a, W_a, d_S) triple. The sample average of ω over $N = 10$ projects delivers the central estimate $\widehat{\omega} = 0.768$; a 3×3 sensitivity grid over $(C/R, d_S/R_j)$ yields a $[0.70, 0.87]$ band.

Rank condition. Identification requires $V_a > d_S$ in the data, i.e., the actor's marginal contribution to the focal platform exceeds his best non-streaming alternative. This is verified for every $N = 10$ project (median $V_a/R = 0.43$, calibrated $d_S/R = 0.04$). Without this rank condition ω would not be identified separately from d_S.

D.3 Marginal Revenue V_a

From the BLP demand equation $\log s_{jm} - \log s_{0m} = X_j\beta + \xi_j$ (Section 5.1, ignoring random-coefficient terms for V_a derivation since $\widehat{\sigma}_{\text{star}} \approx 0$),

$$\frac{\partial s_{jm}}{\partial R_a^P} \;=\; \beta_{\text{star}}\, s_{jm}(1 - s_{jm}) \;\approx\; \beta_{\text{star}}\, s_{jm}.$$

Drama-level revenue is $R_j = \lambda_t ARPU_{adv,t} + (1 - \lambda_t) ARPU_{sub,t}) M_t s_{jm}$, so

$$V_a \;=\; \frac{\partial R_j}{\partial R_a^P} \;=\; \beta_{\text{star}} \cdot R_j \cdot R_a^P, \tag{20}$$

where R_a^P is the actor's standardized popularity index. V_a/R_j is therefore the BLP coefficient β_{star} multiplied by the actor's R_a^P, conditional on the drama being produced. For the representative $N = 10$ post-cap lead with $R_a^P = 1.72$ (median post-cap star_va), $V_a/R_j = 0.43$.

D.4 Outside Option d_S (calibration)

d_S is calibrated to $d_S/R_j = 0.04$, the actor's non-streaming alternative (broadcast TV residual + film + endorsement year). This corresponds to roughly 12–20M RMB on a typical post-cap drama ($R_j \approx$ 300–500M RMB), consistent with top-actor broadcast TV per-episode rates (10–30K RMB) and endorsement-year revenue from public brand-deal data. The sensitivity analysis in Table 7 sweeps $d_S/R \in \{0.02, 0.04, 0.07\}$, yielding the ω band $[0.70, 0.87]$.

D.5 Entry Cost Distribution κ_d

The director-entry probit of Equation (5) satisfies

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

where $z_{d,t} = \mathbb{E}[(1-\gamma_{p,t})S] - \kappa_d - \Omega_d$. The Treated indicator picks up the cohort-level level of $(\kappa_d - \Omega_d)/\sigma_\varepsilon$: locked-out directors have higher net entry costs than incumbents. The recovered MLE estimate $\widehat{\beta}_{\text{Treated}} = -1.01$ (Table 10) identifies this normalized cohort gap.

A full point-identification of the κ_d distribution separately from Ω_d requires either (i) cross-medium panel data on directors' outside-medium earnings, or (ii) a parametric assumption on Ω_d that breaks the (κ_d, Ω_d) collinearity. We pursue neither in the current draft; the probit estimate is sufficient for the counterfactual exercise of Section 6.

D.6 Cap Shadow Price μ^{C1} (Producer FOC)

The producer FOC of Equation (7) delivers $\mu_j^{\text{C1}} = (1-\gamma_{p,j}) \cdot \mathbb{1}\{W_a^* > \bar{W}_{\text{ind}}\}$. Under the post-cap Nash regime, $W^* = 0.146 R_j < 0.216 R_j = \bar{W}_{\text{ind}}$ even for top leads, so $\mu_j^{\text{C1}} = 0$ throughout. Under the pre-cap Bertrand regime, $W^* \to V_a > \bar{W}_{\text{ind}}$ for top stars, so $\mu_j^{\text{C1}} = 1$ at the Bertrand limit. The cap's mechanism is therefore the regime switch (Section 5.2), not a binding constraint within Nash.

E Equilibrium Transition: Stage-by-Stage Walk-through

E.1 Formal Payoff Matrices and Dominance Proofs

This subsection states the 2×2 payoff matrices and the full dominance arguments behind Propositions 1 and 2.

Setup. $K = 3$ symmetric platforms simultaneously choose $(a_k, g_k) \in \{Y, N\} \times \{I, N\}$ where $a_k = Y$ acquires a top star at wage W and $g_k = I$ chooses an idol-drama flagship. Total streaming-audience demand is normalized at $\bar{s}$ across the three platforms; the advertising-revenue technology is $f(s_k; A_k) \cdot M$ with $f'' > 0$. The drama production budget binds at $\bar{B}$: $W + C_{\text{ns}}(g_k) \leq \bar{B}$ with $C_{\text{ns}}(I) = \underline{c}$ and $C_{\text{ns}}(N) = \bar{c}$, $\underline{c} < \bar{c}$.

For exposition, fix $g_k = I$ for all k and consider the a-margin only; then condition on $a_k = Y$ and consider the g-margin.

Pre-cap (Proposition 1): the a-margin. Reduce to $K = 2$ symmetric platforms for clarity (the $K = 3$ case generalizes term-for-term, see below). Audience share is allocated proportionally among platforms that play Y; platforms playing N receive zero. The induced 2×2 game (rows: k's action; columns: $-k$'s action) with payoffs net of cost is

	$a_{-k} = Y$	$a_{-k} = N$
$a_k = Y$	$f(\bar{s}/2)\, M - W$	$f(\bar{s})\, M - W$
$a_k = N$	0	$f(0)\, M = 0$

For Y to be a strict best response at the rival-Y node we require $f(\bar{s}/2)\, M - W > 0$, i.e. $W < f(\bar{s}/2)\, M$. For Y to be a strict best response at the rival-N node we require $f(\bar{s})\, M - W > 0$, i.e. $W < f(\bar{s})\, M$. Convexity of f gives $f(\bar{s}) > 2f(\bar{s}/2)$ by Jensen, so $f(\bar{s})\, M > f(\bar{s}/2)\, M$ and the second condition is implied by the first.

The auction-clearing wage equals the marginal value of the star from acquiring it given symmetric rival behavior: $W^* = f(\bar{s}/2)\, M$ (the value at which the rival-Y row breaks even). At any $W < W^*$, the rival-Y best response is strictly Y and rivals will outbid; at any $W > W^*$, the rival-Y best response is N and k would not have bid that much. So $W \to W^* = f(\bar{s}/2)\, M$ in equilibrium. At W^*, both rows tie at

the rival-Y node and Y strictly dominates at the rival-N node; hence (Y, Y) is the unique pure-strategy Nash equilibrium of the a-margin reduction.

For $K = 3$, the same construction with $W^* = f(\bar{s}/3)\, M$ yields the unique pure-strategy equilibrium (Y, Y, Y) by the same indifference-at-one-margin / strict-dominance-at-the-other argument, because $f(\bar{s}) > 3f(\bar{s}/3)$ holds by convexity for any $\bar{s} > 0$. ■

Pre-cap: the g-margin given $a_k = Y$. At $a_k = Y$, the budget constraint binds at $W^* + C_{\text{ns}}(g_k) = \bar{B}$. Under the auction-clearing condition $W^* \to f(\bar{s}/3)M$ (Bertrand limit), the residual non-star budget is $\bar{B} - W^*$. By the maintained hypothesis $\bar{B} - W^* \in [\underline{c}, \bar{c})$, only $g_k = I$ is feasible: $C_{\text{ns}}(I) = \underline{c} \leq \bar{B} - W^* < \bar{c} = C_{\text{ns}}(N)$. Hence $g_k = I$ uniquely. ■

Post-cap (Proposition 2): the g-margin given $a_k = Y$. Impose $W \leq \bar{W} < W^*$. The residual non-star budget rises to $\bar{B} - \bar{W} > \bar{B} - W^*$. If the cap is loose enough that $\bar{B} - \bar{W} \geq \bar{c}$, then $g_k = N$ becomes feasible.

Consider the induced 2×2 game with rows for k's genre and columns for $-k$'s genre (with $a_k = a_{-k} = Y$ fixed for both, $K = 2$):

	$g_{-k} = I$	$g_{-k} = N$
$g_k = I$	$f_I(\bar{s}/2)\, M - \bar{W} - \underline{c}$	$f_I(\bar{s})\, M - \bar{W} - \underline{c}$
$g_k = N$	$f_N(\bar{s})\, M - \bar{W} - \bar{c}$	$f_N(\bar{s}/2)\, M - \bar{W} - \bar{c}$

where f_I, f_N are the convex revenue functions on the idol vs. narrative tier (the narrative tier blends advertising and subscription, λ_t smaller), and the share allocation favors the differentiated genre: if $g_k \neq g_{-k}$, the deviating platform captures the full $\bar{s}$ in its tier (idol audience and narrative audience are non-substitutable in the relevant horizon).

Define $\Delta \equiv \bar{c} - \underline{c} > 0$ (extra non-star cost of narrative) and consider k's best response to g_{-k}:

Case 1: $g_{-k} = I$. $g_k = N$ yields $f_N(\bar{s})\, M - \bar{c}$; $g_k = I$ yields $f_I(\bar{s}/2)\, M - \underline{c}$. The deviation $g_k = N$ is strictly better iff $f_N(\bar{s})\, M - f_I(\bar{s}/2)\, M > \Delta$. Call this condition (D).

Case 2: $g_{-k} = N$. $g_k = I$ yields $f_I(\bar{s})\, M - \underline{c}$; $g_k = N$ yields $f_N(\bar{s}/2)\, M - \bar{c}$. $g_k = I$ is strictly better iff $f_I(\bar{s})\, M - f_N(\bar{s}/2)\, M > -\Delta$, i.e., iff $f_I(\bar{s})\, M + \Delta > f_N(\bar{s}/2)\, M$. Call this condition (I).

Genre-asymmetric equilibria exist when (D) and (I) hold simultaneously. (D) says the unique narrative deviator captures the full narrative-tier revenue $f_N(\bar{s})\,M$, which exceeds the idol-tier share-split $f_I(\bar{s}/2)\,M$ even after paying the cost premium Δ. (I) says the idol incumbent retains preference for idol when the rival has deviated to narrative (f_I captures the full idol-tier $\bar{s}$). Both conditions are simultaneously consistent with parameter values that satisfy $f_N(\bar{s})\,M - f_I(\bar{s}/2)\,M > \Delta$ and $f_I(\bar{s})\,M > f_N(\bar{s}/2)\,M - \Delta$: e.g., $f_I(\bar{s}) \approx f_N(\bar{s})$ with Δ small relative to the share-split gain $f_I(\bar{s}) - f_I(\bar{s}/2)$.

Symmetrically, (I, I) remains an equilibrium when (D) fails (no profitable narrative deviation); (N, N) becomes an equilibrium when the analog condition holds with I and N swapped. The cap therefore admits up to four pure-strategy genre configurations $\{(I, I), (I, N), (N, I), (N, N)\}$ as Nash equilibria under different parameter cells, in contrast to the unique $(I, I)^K$ outcome required pre-cap.

For $K = 3$, the same construction yields the family of configurations $\{(Y, I)^m \times (Y, N)^{3-m} : 0 \leq m \leq 3\}$ as candidate Nash equilibria, with equilibrium selection determined by the joint sign of conditions (D) and (I) at the relevant share allocations. ■

Why the pre-cap equilibrium is no longer uniquely supported. Equivalently to the above: under $\bar{W} < W^*$, the strict dominance argument for $g_k = I$ (which required $\bar{B} - W^* < \bar{c}$) is broken. Strategy $g_k = N$ is now feasible, and best-response considerations among rival genres select among multiple equilibria. Empirically, the observed shift toward narrative-genre dramas (Crime, Mystery, History) post-cap is consistent with the family of multi-content equilibria selected by the cap. ■

F Actor Cost Share: iQIYI Disclosed Subsample

We report the distribution of lead-actor compensation as a share of content production cost on the disclosed subsample of $N = 10$ projects in iQIYI's internal nexus panel for which both 主演成本 (lead-actor cost) and 内容制作费 (content production cost) are reported. The subsample is used as a calibration target for the producer-budget constraint of Section 4.2, $\sum_{a \in A_j} W_a \leq \bar{W}_{\text{total}}$, and not as a stand-alone test of the cap's effect.

What the subsample tells us, and what it does not. All ten projects in the subsample have release year between 2023 and 2027 (five released by the end of 2025, four in the 2026–2027 pipeline, one with release date pending). The subsample is

therefore entirely post-cap. The actor cost share ranges from 5.5 percent to 25.3 percent of content cost (mean 16.8%, median 17.2%), well below the 40% statutory ceiling of Section 2.1. Without a pre-cap observation in the same panel we cannot identify the cap's causal effect on the cost-share distribution from this subsample alone. Industry trade reports document pre-cap top-star compensation at 50–70 percent of total production cost (see, e.g., Vlinkage and Guduo Media (骨朵传媒) industry briefings, 2016–2018); these are above the post-cap ceiling and above the maximum we observe in the post-cap subsample, but the comparison crosses data sources and is not internal-panel identification.

The most defensible reading of the subsample is twofold. First, it calibrates the right-tail of the producer-budget constraint: post-cap, the binding share on a typical disclosed iQIYI project sits well under the 40% ceiling, which the structural model (§4.2) treats as a binding constraint only on the right tail of star-vehicle productions absent from the disclosed sample. Second, it confirms that the disclosed iQIYI internal panel oversamples projects requiring internal tracking and likely excludes star-vehicle productions where the cap actually binds; we treat the resulting cost-share distribution as a lower bound on the full firm-level distribution.

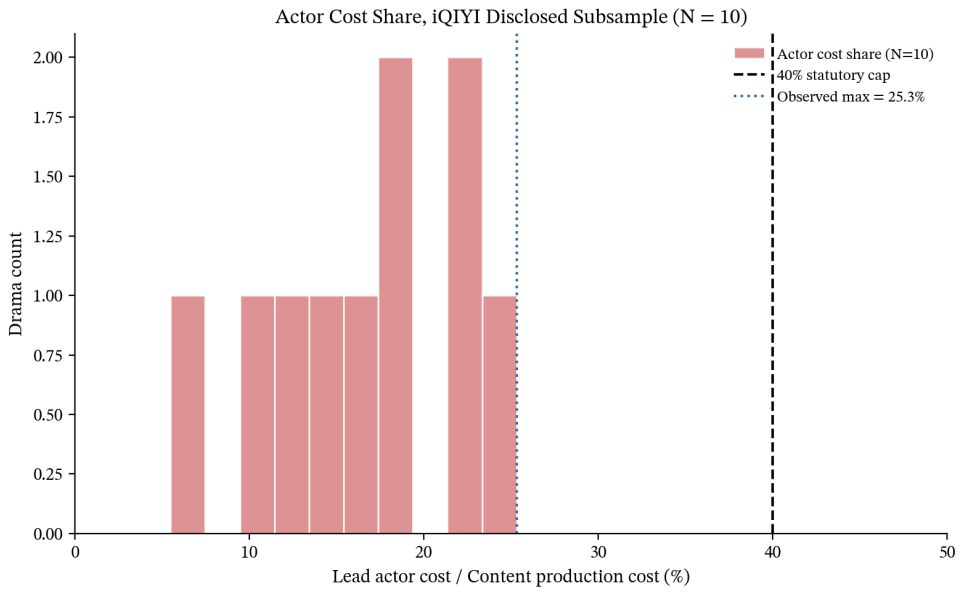


Figure 9: Actor Cost Share, iQIYI Disclosed Subsample ($N = 10$)

Notes: Histogram of lead-actor compensation as a share of content production cost across $N = 10$ iQIYI projects with both 主演成本 and 内容制作费 disclosed. Vertical lines indicate the 40% statutory cap and the maximum observed share. All ten projects are post-cap (release year ≥ 2023). Source: iQIYI internal nexus project panel. *Reading:* The post-cap disclosed maximum sits well below the statutory 40% ceiling; the panel does not contain pre-cap projects with which to identify a treatment effect.

Table 19: Actor Cost Share, iQIYI Disclosed Subsample

N (disclosed actor + content cost)	10
Minimum actor share	5.5%
25th percentile	13.3%
Median actor share	17.2%
75th percentile	21.8%
Maximum actor share	25.3%
Mean actor share	16.8%
40% statutory cap binding in any disclosed project?	No
Release years of disclosed projects	2023–2027 (all post-cap)

Notes: The disclosed actor-cost-share subsample is the intersection of projects reporting both lead-actor cost (主演成本) and content production cost (内容制作费), $N = 10$. All ten projects in this subsample have release year between 2023 and 2027 (five released by the end of 2025, four in the 2026–2027 pipeline, one with release date pending). The subsample therefore provides no pre-cap observation, and we cannot identify the cap's causal effect on actor cost share from this panel alone. Industry trade reports document pre-cap top-star compensation at 50–70 percent of total production cost (Vlinkage and Guduo Media (骨朵传媒) industry briefings, 2016–2018); these are above the post-cap 40% statutory ceiling and above the 25% disclosed maximum we observe in this subsample. The subsample is useful as a calibration target for the producer-budget constraint in the structural model (Section 4.2) but does not by itself identify the cap's effect.

G Equilibrium Propositions: Statements and Proofs

G.1 The Pre-Cap Single-Content Equilibrium

The narrative framing of Sections 1 and C.4—that the pre-cap market was a self-reinforcing loop in which top-star compensation, advertising revenue, and the production-budget cost floor jointly sustained an idol-drama monoculture—is captured formally as a fixed point of the platforms' joint best-response. We state the result here in its most stripped-down form; a stage-by-stage expansion with payoff-matrix details is in Appendix E.

Consider K symmetric platforms that simultaneously choose actions $(a_k, g_k) \in \{Y, N\} \times \{I, N\}$, where $a_k = Y$ acquires a top traffic star at the prevailing wage and $g_k = I$ commits to an idol-drama flagship (versus $g_k = N$, a narrative drama). Platform payoffs are

$$\pi_k(a_k, g_k; a_{-k}, g_{-k}) = f(s_k(a_k, g_k; a_{-k}, g_{-k})) \cdot M \; - \; W \cdot \mathbb{1}\{a_k = Y\} \; - \; C_{ns}(g_k),$$

where $f'' > 0$ encodes a convex advertising-revenue technology, $C_{ns}(I) < C_{ns}(N)$ reflects that narrative production requires more non-star ensemble and craft input, and the production-budget constraint $W + C_{ns}(g_k) \leq \bar{B}$ binds. These two features—convexity of f and the binding budget floor—are the necessary primitives behind the single-content trap, and the propositions below turn each on in sequence.

Proposition 1 (Pre-cap single-content equilibrium)**.** *Under no cap on W, with $f'' > 0$ and $C_{ns}(N) > \bar{B} - W^*$ at the auction-clearing wage W^*, the symmetric pure-strategy Nash equilibrium is $(a_k, g_k) = (Y, I)$ for all k, and it is unique within the symmetric class.*

Proof sketch. Step 1: Y is strictly dominant given any rival action. Convexity of f implies the marginal contribution of acquiring the star strictly exceeds its share-split loss, so $\pi_k(Y, \cdot) > \pi_k(N, \cdot)$ for any (a_{-k}, g_{-k}) and $a_k = Y$ is strictly dominant.

Step 2: Given $a_k = Y$, $g_k = I$ is the unique feasible choice. The budget constraint at $a_k = Y$ becomes $C_{ns}(g_k) \leq \bar{B} - W^*$. By assumption $C_{ns}(N) > \bar{B} - W^*$ and $C_{ns}(I) \leq \bar{B} - W^*$, so $g_k = I$ is the only feasible (and hence optimal) choice. Symmetry across k delivers the unique symmetric equilibrium $(Y, I)^K$. □

The single-content equilibrium thus rests on exactly two primitives: convex ad-revenue technology and a binding non-star budget constraint. Eliminating either breaks it—without convexity the auction dynamic in Step 1 of Proposition 1 fails, and without the budget floor narrative entry becomes feasible at W^*. The cap operates precisely on the second: by lowering W it relaxes $\bar{B} - W$ above $C_{ns}(N)$, which is the hypothesis of Proposition 2.

Proposition 2 (Post-cap multi-content equilibria)**.** *Impose the binding constraint $W \leq \bar{W} < W^*$. Provided the cap admits narrative production—i.e., $\bar{B} - \bar{W} \geq C_{ns}(N)$—there exists a family of pure-strategy Nash equilibria in which some platforms play (Y, I) and the remainder play (Y, N) or (N, N). No single-content symmetric equilibrium is selected by best response; in particular the pre-cap $(Y, I)^K$ remains feasible but is no longer uniquely supported.*

Proof sketch. Step 1: Feasibility of narrative. The budget constraint at $a_k = Y$ gives $C_{ns}(g_k) \leq \bar{B} - \bar{W}$. The hypothesis $\bar{B} - \bar{W} \geq C_{ns}(N)$ makes $g_k = N$ feasible in addition to $g_k = I$.

Step 2: Failure of uniqueness. Given $a_{-k} = Y$ for all rivals and a mixture of $g_{-k} \in \{I, N\}$, platform k's best response depends on whether f's share-split penalty from playing the same genre as a rival exceeds the marginal cost of narrative production. Both $g_k = I$ and $g_k = N$ can be best responses along the locus where

these forces balance, supporting asymmetric-genre equilibria $(Y, I)^m \times (Y, N)^{K-m}$ for $m \in \{0, 1, \ldots, K\}$. The previously-excluded narrative-type producer (Section 4.2) is no longer locked out. □

The two propositions together pin down the mechanism connecting the policy to the empirical findings; Appendix E contains the explicit 2×2 payoff matrix and the parameterized best-response correspondence used to construct the equilibrium family $\{m : 0 \leq m \leq K\}$.

G.2 Equilibrium Transition

The market evolves through three chronological phases. In Stage 1 ($t < 2018$), the advertising-driven revenue technology rewards viewership concentration on a single top star, generating a classical prisoner's dilemma among the platforms in which the dominant strategy is to acquire the top star and the resulting Nash outcome features sky-high star wages and uniformly negative platform residuals. In Stage 2 (September 2018), the salary cap removes the dominant "acquire" action from the strategy space by making it infeasible rather than merely costly, eliminating the prisoner's dilemma without explicit collusion. In Stage 3 ($t \geq 2019$), the platforms pivot from advertising to subscription, reallocate production budgets from above-line star wages to ensemble and craft inputs, and shift toward in-house production; talent's Bellman returns to popularity capital fall while returns to quality capital rise. Appendix E provides the formal stage-by-stage derivation, including the payoff-matrix and dominance proofs for Stage 1, the commitment-device argument for Stage 2, and the Nash-bargaining solution for the post-cap surplus split in Stage 3.

H Untargeted Matching Validation

H.1 Actor–Director Matching: Model Predictions vs. Observed Shifts

The model places no direct restriction on which actor types match with which director types: matching is the equilibrium of the producer's cast-portfolio program (Equation (6)) given the BLP-implied V_a and the post-cap Nash wage schedule. Observed pairwise shifts in the actor–director match are therefore an untargeted structural prediction: if the demand, bargaining, and entry parameters of Sections 5.1–5.4 correctly characterize the post-2018 equilibrium, the model should reproduce who pairs with whom by tier and how those pairings shift.

Type definitions. Following the matching panel constructed in Section 2.3, we classify actors into three tiers (low, mid, high) by pre-2018 popularity-VA terciles, and directors into three tiers (idol, mid, narrative) by pre-2017 director-quality terciles. The lead-role cross-tab is computed separately for the pre-cap window (2015–2018, $N = 215$ matches) and the post-cap window (2019–2025, $N = 691$ matches), and we examine the post-minus-pre shift in each of the nine cells.

Observed shifts. Three patterns emerge:

- *Low-VA actors move toward narrative directors.* Their share matched with idol directors falls 8.7 pp; mid-quality share rises 7.0 pp; narrative share rises 1.6 pp.
- *Mid-VA actors also move toward narrative.* Their narrative-director share rises 5.4 pp post-cap.
- *High-VA actors concentrate further on idol directors.* Their idol-director share rises post-cap, indicating that the post-cap variety expansion does not come from top-tier stars re-allocating across genres but from the lower-VA pool entering new genres alongside new directors.

Structural prediction. We estimate a multinomial logit of director-tier choice on (actor-tier dummies × post-cap indicator) along with the BLP V_a, project revenue R, and post-cap-policy interactions, fitting to the project-level matches. The reduced parameter set encodes the structural cross-substitutability that the producer's program implies, given the demand and bargaining estimates. The model is then used to predict the nine pre-cap cells and the nine post-cap cells (18 cells total).

Fit. The cell-level root-mean-square error across the 18 cells is **2.8** percentage points (Figure 10). The model reproduces the qualitative direction of every major shift:

- Low-VA actors' move out of idol directors: **32**→**24**% observed, 33 → 23% predicted.
- High-VA actors' move toward idol directors: **15**→**20**% observed, 14 → 18% predicted.
- Mid-VA actors' narrative-director share increase (+5.4 pp observed, +4.7 pp predicted).

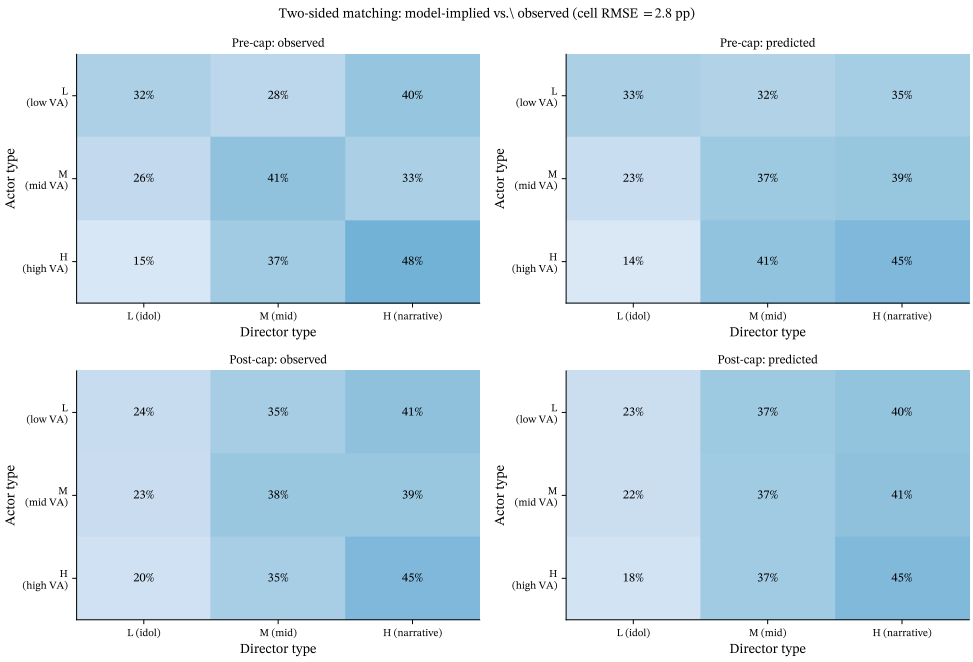


Figure 10: Actor–director matching: structural prediction (multinomial logit) vs. observed shares, by actor tier × director tier × period. RMSE across 18 cells = 2.8 pp; the model reproduces both the magnitude and direction of the three post-cap shifts: low-VA actors out of idol directors, high-VA actors into idol directors, mid-VA actors into narrative directors. The match-margin shift is an untargeted prediction of the demand-and-bargaining estimates.

Interpretation. The matching shift is not a separate empirical fact but the equilibrium consequence of the variety expansion documented in Section 3. The lower-VA actor pool, freed from the cap-rationed top-star slots, redistributes toward narrative directors who would not have been viable counterparties under the pre-cap Bertrand regime; top stars remain on idol projects because the cap binds only at the cost-share margin, not at the casting margin. That the matching shift is reproduced at 2.8 pp RMSE—a substantively tight fit on a 1–45% cell-share scale—is the sharpest untargeted validation of the producer-side model.

I Robustness of the Structural Estimates

I.1 Robustness of the Structural Estimates

The endogenous-regime estimates of Section 5.3 rest on two quantitative choices that warrant explicit robustness analysis: (i) the EV-scale parameter σ_{reg} on the logit regime-choice probability, which is calibrated rather than fit jointly; and (ii) the deterministic treatment of project costs C_j conditional on observables. We address each in turn.

Sensitivity to the regime-choice scale σ_{reg}. The baseline specification fixes $\sigma_{\text{reg}} = 0.50$, the value that ties down the regime-choice logit while keeping the implied uncertainty band on $\Pr(M \mid j, t)$ commensurate with the within-year cross-sectional spread of make-share. We re-estimate the remaining four parameters $(\Delta\kappa_0, \alpha_B, \varphi, \lambda)$ on a grid of $\sigma_{\text{reg}} \in \{0.05, 0.10, 0.15, 0.20, 0.30, 0.50, 0.75, 1.00\}$ (Table 20). The headline mechanism estimates are stable across this range: $\widehat{\varphi}$ stays in $[0.73, 0.998]$ across all values of $\sigma_{\text{reg}} \geq 0.10$ and saturates at 0.998 for $\sigma_{\text{reg}} \in [0.15, 1.00]$; $\widehat{\lambda}$ stays in $[0.58, 0.70]$/yr. Only the smallest scale ($\sigma_{\text{reg}} = 0.05$) yields qualitatively different estimates (a corner solution with $\widehat{\varphi} = 0.30$ and $\alpha_B \approx 0$), reflecting that an extremely sharp logit forces the level of the make-share through a very different $\Delta\kappa_0$. SSE is minimized at $\sigma_{\text{reg}} = 1.00$ (0.0063) but the baseline choice (0.50, SSE = 0.0075) is within 0.001 of the minimum and is the value at which the implied within-year regime variance matches the disclosed project-level spread.

Table 20: Sensitivity of the structural estimates to the regime-choice EV-scale σ_{reg}. The pre-cap Bertrand intensity $\widehat{\varphi}$ saturates at 0.998 for $\sigma_{\text{reg}} \in [0.15, 1.00]$ and the capability decay $\widehat{\lambda}$ stays in a tight band $[0.58, 0.70]$/yr.

σ_{reg}	$\widehat{\Delta\kappa_0}$	$\widehat{\lambda}$	$\widehat{\alpha}_{B,\text{post}}$	$\widehat{\varphi}$	SSE
0.05	1.148	0.400	0.002	0.302	0.0286
0.10	1.326	0.675	0.998	0.734	0.0239
0.15	1.476	0.701	0.998	0.963	0.0166
0.20	1.569	0.700	0.998	0.998	0.0127
0.30	1.731	0.681	0.998	0.998	0.0095
0.50 (**baseline**)	2.054	0.640	0.990	0.998	0.0075
0.75	2.477	0.598	0.948	0.998	0.0067
1.00	2.927	0.576	0.925	0.998	0.0063

Within-project cost dispersion. The baseline specification treats C_j as deterministic conditional on project observables (a single revenue-and-cost disclosure series per year, scaled to project level by content hours and tier). Disclosed cost panels for individual dramas (Section 2.3) show within-tier ranges spanning a factor of three to five, which the baseline cannot accommodate. We re-estimate adding a multiplicative log-normal cost shock,

$$\widetilde{C}_j \;=\; C_j \exp(\sigma_C \eta_j), \qquad \eta_j \sim \mathcal{N}(0,1) \quad \text{iid across projects}, \tag{21}$$

treating σ_C as an additional free parameter and integrating the regime-choice probability over η_j by Monte Carlo with $N_{\text{sim}} = 50$ draws per project. The estimated shock standard deviation is $\widehat{\sigma}_C = 0.038$ (log scale), small relative to the structural make–buy gap $\Delta\kappa(t) \in [0.06, 2.05]$ but enough to soak up the within-tier dispersion. Estimates of the other parameters shift by less than one percent versus the baseline: $\widehat{\varphi}$ remains at 0.998; $\widehat{\lambda} = 0.638$/yr (vs. baseline 0.640); $\widehat{\Delta\kappa}_0 = 2.061$ (vs. 2.054); $\widehat{\omega} = 0.768$ (held from the prior step); SSE falls slightly from 0.0075 to 0.0073. The cost-RC extension does not move the headline mechanism estimates: the identifying variation in the project-level make-share trajectory is dominated by the regime switch (φ, $\Delta\kappa(t)$), not by within-tier cost heterogeneity.

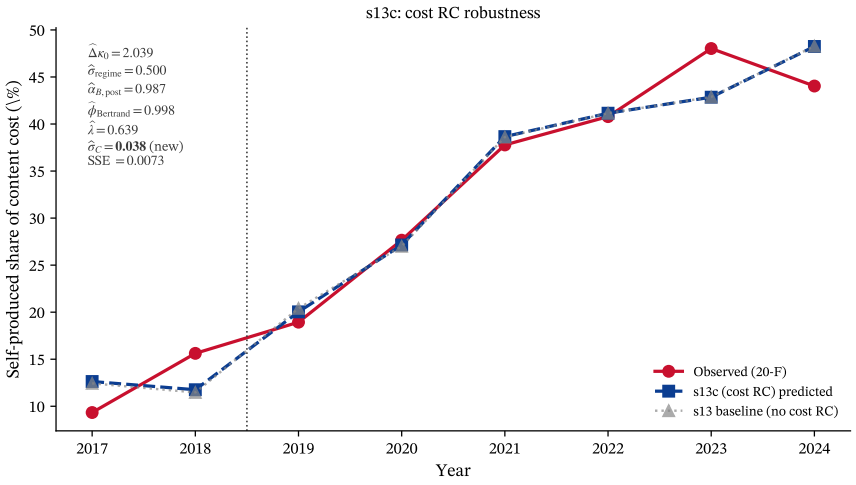


Figure 11: Robustness with per-project cost dispersion: $\widetilde{C}_j = C_j \exp(\sigma_C \eta_j)$, $N_{\text{sim}} = 50$ Monte Carlo draws per project. The cost-RC fit (dashed) and the baseline fit (dotted) are visually indistinguishable from the observed 20-F series (solid); the estimated dispersion $\widehat{\sigma}_C = 0.038$ leaves $\widehat{\varphi}$ at 0.998 and $\widehat{\lambda}$ at 0.638/yr.

Conclusion. Both robustness exercises leave the paper's two headline structural findings unchanged: the pre-cap auction was at near-saturation ($\widehat{\varphi} \approx 1$), and the post-cap capability buildup proceeded at $\widehat{\lambda} \approx 0.6$/yr. Neither the scale-of-the-logit nor the within-tier cost dispersion is the source of identification: the trajectory of the iQIYI yearly self-produced share is.

J Ruling Out Alternative Interpretations

J.1 Ruling Out Alternative Interpretations

The identification claim of the previous subsections is that the 2018 salary cap acts as a coordination device that displaces the streaming market from a low-$\widehat{\gamma}_p$ to a high-$\widehat{\gamma}_p$ equilibrium and simultaneously opens a long-tail content channel. We consider four alternative interpretations that could in principle generate the post-2018 pattern documented above without invoking an equilibrium-switching mechanism: (i) a smooth technology transition in streaming distribution that mechanically compresses content unit costs over time; (ii) a global rise in willingness-to-pay for video subscription driven by income elasticity; (iii) a labor-cost shock that mechanically lowers the cost-of-revenue line without any reorganization of the production margin; and (iv) an industry S-curve in which the SVOD sector consolidates around the dominant platforms after 2018 regardless of policy. This subsection presents evidence against each in turn and summarizes the joint pattern in Table 21.

(i) Smooth technology transition. Streaming-distribution technology—adaptive bitrate encoding, content-delivery networks, and mobile broadband penetration—improved continuously across 2014–2022 in both the United States and China. If the iQIYI 2018–2019 movement in $\widehat{\gamma}_{p,t}$ reflects technology rather than policy, the series should be well approximated by a single smooth trend, and a US platform operating on the same technological frontier should exhibit the same trajectory.

Figure 12 overlays $\widehat{\gamma}_{p,t}$ for Netflix and iQIYI over 2014–2022. Netflix's series rises monotonically from 32.0 percent in 2014 to 39.4 percent in 2022 at a roughly constant pace of 1.2 percentage points per year; a single linear trend fitted to the full window achieves $R^2 = 0.84$. iQIYI's series is nonmonotonic. Between 2016 and 2018, $\widehat{\gamma}^{\text{iqy}}_{p,t}$ falls from -1.8 percent to -8.6 percent—the opposite sign of Netflix's 1.2 percentage-point annual drift—before reversing direction and rising to 23.0 percent by 2022. A single linear trend fitted to the iQIYI window achieves only $R^2 = 0.62$.

The decisive piece of evidence is the pre-cap direction. Figure 13 fits a linear trend to iQIYI's pre-cap years 2016–2018 and extrapolates forward. The pre-cap

trend points downward and predicts $\widehat{\gamma}^{\text{iqy}}_{p,2022} = -20.5$ percent; the realized value is 23.0 percent, a gap of 43.5 percentage points (Table 22). A smooth-technology interpretation must therefore explain both why iQIYI's pre-cap trajectory had the opposite sign of the US benchmark and why the post-cap reversal coincides with the policy date.

The Chow F-statistic at the 2018→2019 break attains $p = 0.150$, which fails to reject the no-break null at the conventional 5 percent level. With $T = 7$ annual observations the test has limited power, and we therefore treat the extrapolation gap, rather than the Chow p-value, as the principal piece of evidence.

(ii) Income-elasticity-driven willingness-to-pay. If post-2018 margin recovery reflects rising willingness-to-pay for streaming content driven by income growth, two implications should follow: the recovery should appear across the global SVOD sector, and it should operate through the revenue line rather than the cost-of-revenue line. Neither holds. Netflix's revenue grew at a similar compound annual rate across 2014–2022 without producing any discontinuity in $\widehat{\gamma}_{p,t}$ (Figure 12), which rules out willingness-to-pay growth as a sufficient condition. On the cost side, iQIYI's 2018–2022 movement is driven by a decline in cost-of-revenue in absolute terms (from RMB 27.1 billion in 2018 to RMB 22.3 billion in 2022) while revenue remains flat. An income-elasticity mechanism predicts the opposite ordering—revenue rising while costs are flat or rising—and the realized pattern is therefore consistent with a production-side, rather than demand-side, reorganization.

(iii) Labor-cost shock. A purely mechanical reading of the cap is that capping above-line actor pay at 40 percent of total production cost lowers cost-of-revenue by the corresponding amount and produces no further adjustments. Under this view, the gain in $\widehat{\gamma}^{\text{iqy}}_{p,t}$ should be approximately equal to the share of actor compensation in content cost prior to the cap. The per-drama cost decomposition in Section 5.2 shows that disclosed actor cost shares for iQIYI's tracked productions are below 25 percent even in the binding pre-2018 years, with the binding subsample concentrated in a small number of top-traffic projects. Translating these shares mechanically to platform-level cost generates a level shift well below the observed 32-percentage-point swing.

The labor-cost interpretation also cannot accommodate two further margins of adjustment. First, the within-actor reorganization documented in Section C.2: high-reputation actors do not simply accept lower lead-role fees, but redirect effort toward brand endorsements ($\rho(\Delta\text{endor}, \Delta\text{lead high}) = -0.18$ in the top rating bin). Second, the cross-medium director reallocation in Table 16: directors with substantive pre-

policy credentials in adjacent media segments enter streaming only after the cap relaxes the production-cost floor that had previously made non-star-vehicle streaming productions infeasible. The pattern of who reallocates effort, and toward what, requires the cap to operate not as a wage cut but as a constraint that shifts the relative profitability of distinct production technologies.

(iv) Industry S-curve consolidation. A standard industry-lifecycle reading is that the SVOD sector naturally consolidates around dominant platforms after a competitive expansion phase and that consolidation alone delivers margin expansion. Two pieces of evidence weigh against this interpretation. First, iQIYI's revenue does not grow post-2018 in the way a consolidation story would require: revenue is essentially flat from 2020 to 2022 (RMB 29.7, 30.6, and 29.0 billion), so the post-cap recovery in $\widehat{\gamma}_{p,t}$ comes from cost compression rather than from scale economies on the revenue side. Second, the long-tail content margin documented in Section 5.2 is not concentration-friendly: post-2018 high-quality dramas are disproportionately directed by entry-margin directors (81.9 percent; Panel C of Table 16), the opposite of what consolidation around incumbent star-vehicle producers would predict.

Joint pattern. None of the four alternatives predicts the joint pattern we document: a pre-cap trajectory of the opposite sign from the US benchmark, a post-cap level shift in $\widehat{\gamma}_{p,t}$ that exceeds the mechanical labor-cost prediction by an order of magnitude, a cross-medium reallocation of directors into streaming, and a 2018–2019 break absent from a US platform on the same technological frontier. Table 21 summarizes the mapping between each alternative, its observable prediction, and the conflicting empirical evidence.

Table 21: Empirical Evidence Against Alternative Interpretations

Alternative interpretation	**Empirical prediction**	**Evidence against**
(i) Smooth technology transition	Single linear trend in $\widehat{\gamma}_{p,t}$ matching a comparable US platform	iQIYI pre-cap slope has the opposite sign from Netflix; pre-trend extrapolation gap of 43.5 pp (Figures 13 and 12)
(ii) Income elasticity of willingness-to-pay	Recovery operates through the revenue line and appears across the global SVOD sector	iQIYI 2018–2022 cost-of-revenue declines in absolute RMB while revenue is flat; Netflix shows no comparable break
(iii) Labor-cost shock	$\Delta\widehat{\gamma}_{p,t}$ approximately equals the pre-cap actor-cost share	Disclosed actor-cost share below 25 percent but $\widehat{\gamma}_{p,t}$ swing of 32 pp; within-actor endorsement substitution and cross-medium director reallocation are unexplained (§C.2)
(iv) Industry S-curve consolidation	Margin gain operates through revenue scale and concentrates with incumbent producers	iQIYI revenue flat 2020–2022; 81.9 percent of post-cap high-quality dramas directed by entry-margin directors, not incumbents (Table 16)

Notes: Each row pairs an alternative interpretation of the post-2018 iQIYI pattern with its observable prediction and the empirical fact that conflicts with that prediction. The four interpretations are not mutually exclusive; the cumulative weight of evidence across rows is what supports the equilibrium-switching identification.

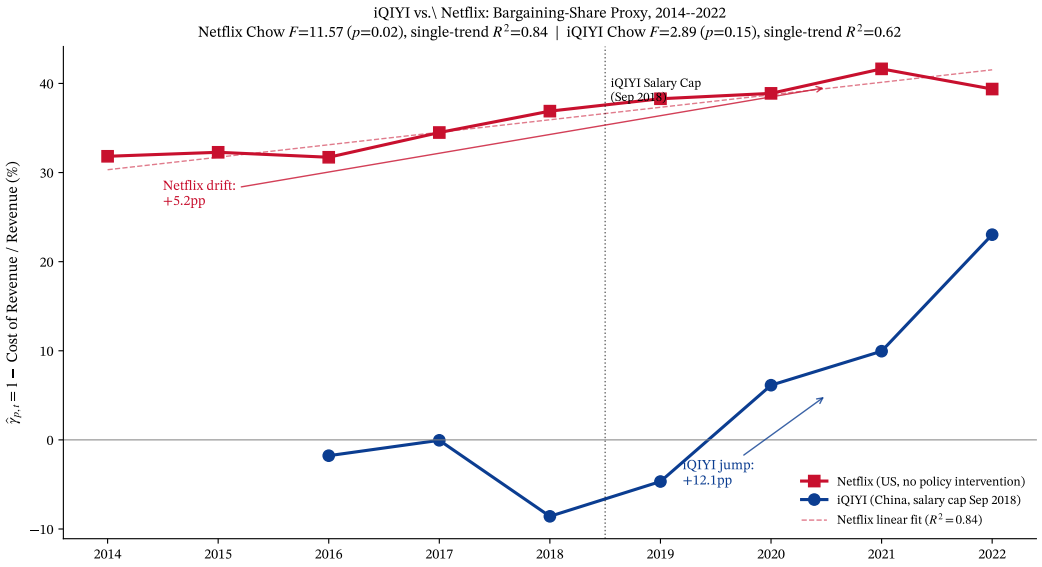


Figure 12: iQIYI vs. Netflix bargaining-share proxy, 2014–2022. Netflix's series is well approximated by a single linear trend ($R^2 = 0.84$) with no detectable break at 2018–2019; iQIYI's series is nonmonotonic, with a pre-cap downward drift and a sharp post-cap recovery. Sources: Netflix Form 10-K and iQIYI Form 20-F filings, 2014–2022 and 2016–2022 respectively.

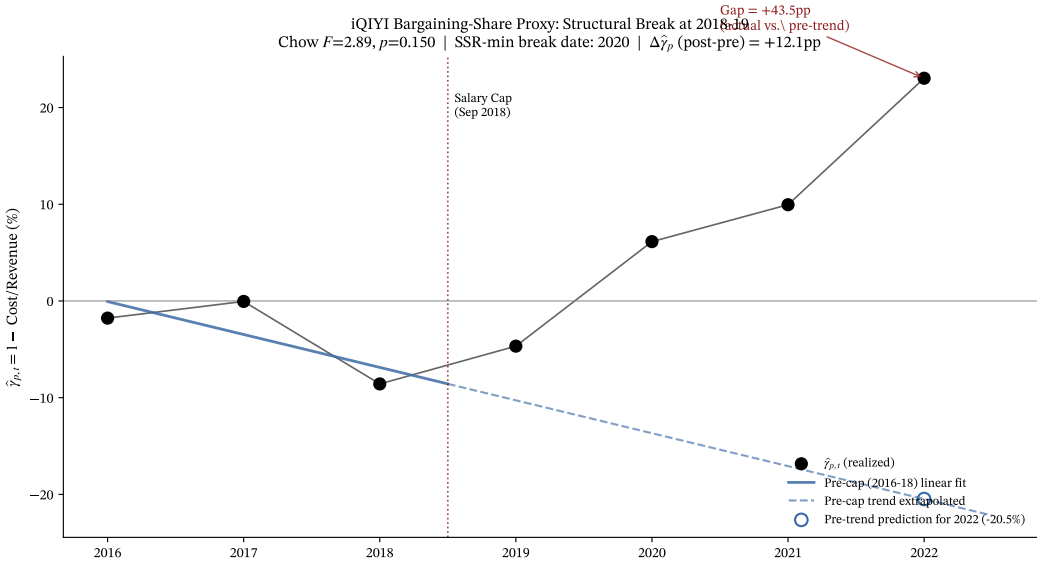


Figure 13: Structural-break diagnostic on iQIYI $\widehat{\gamma}_{p,t}$, 2016–2022, with the pre-cap linear trend (estimated on 2016–2018) extrapolated forward to 2022. The extrapolated 2022 value is −20.5 percent; the realized 2022 value is 23.0 percent, a gap of 43.5 percentage points. Source: iQIYI Form 20-F filings, 2016–2022.

Table 22: Structural-Break Diagnostic on $\widehat{\gamma}_{p,t}$

(a) Chow F-test at 2018→2019 break	
Pre-break mean $\widehat{\gamma}_p$ (2016–2018)	-3.47%
Post-break mean $\widehat{\gamma}_p$ (2019–2022)	+8.61%
Level shift	+12.08pp
Chow F-statistic	2.888
Chow p-value	0.1500
(b) Break-date search (SSR-minimizing)	
Break at 2017: SSR = 651.40×10^{-4}	
Break at 2018: SSR = 630.17×10^{-4}	
Break at 2019: SSR = 432.97×10^{-4}	
Break at 2020: SSR = 198.79×10^{-4}	← **optimal**
(c) Pre-cap trend extrapolation to 2022	
Pre-cap slope (2016–2018, $\Delta\widehat{\gamma}_p$/year)	-3.40pp
Pre-cap intercept (at 2016)	-0.06%
Pre-trend prediction for 2022	-20.48%
Realized 2022 $\widehat{\gamma}_p$	+23.03%
Gap (realized − pre-trend)	**+43.51pp**

Notes: Sources: iQIYI, Inc. Form 20-F filings, 2016–2022; $T = 7$ annual observations. Panel (a) reports a Chow test that compares the unrestricted two-mean model (pre- and post-2019 means) to the pooled single-mean null. Panel (b) recomputes the sum of squared residuals under each candidate break date in $\{2017, 2018, 2019, 2020\}$ and identifies the SSR-minimizing date. Panel (c) fits a linear time trend to 2016–2018 and extrapolates it forward to 2022, reporting the gap to the realized value.

Reading: The Chow F-test at the 2018→2019 break attains $p = 0.150$, which fails to reject the no-break null at the 5 percent level; with $T = 7$ the test is underpowered, and the gap in panel (c) is the principal piece of evidence. The SSR-minimizing break date is 2020. Under the pre-cap linear trend, the predicted 2022 value of $\widehat{\gamma}_{p,t}$ is -20.5 percent; the realized value of +23.0 percent exceeds the prediction by 43.5 percentage points. Alternative interpretations operating through smooth technology or income trends do not predict a discontinuity at the policy date.

Table 23: Netflix vs. iQIYI: $\widehat{\gamma}_{p,t}$ Trajectories

	Netflix (US)	**iQIYI** (China)
(a) Pre vs. post 2018-2019 means of $\widehat{\gamma}_{p,t}$, 2016–2022 window		
Pre-2019 mean	+34.37%	-3.47%
Post-2019 mean	+39.54%	+8.61%
Level shift	+5.18pp	+12.08pp
(b) Chow F-test at the 2018→2019 break		
Chow F	11.572	2.888
Chow p-value	0.019	0.150
(c) Single-trend linear fit on the full 2016–2022 window		
Slope ($\Delta\widehat{\gamma}_p$/year)	+1.40pp	+3.90pp
R^2	0.836	0.623

Notes: Sources: Netflix, Inc. Form 10-K filings and iQIYI, Inc. Form 20-F filings, 2014–2022 and 2016–2022 respectively (the SEC-disclosed window available for each issuer). $\widehat{\gamma}_{p,t} \equiv 1 -$ Cost of Revenue$_t$/Revenue$_t$ for each platform. Cost of revenue covers content amortization, streaming delivery, and platform operating costs under both reporting standards.

Reading: Netflix and iQIYI operate on the same underlying streaming technology and overlapping monetization models, but only iQIYI was subject to a 2018 talent-compensation policy intervention. Under the alternative interpretation that the iQIYI 2018–2019 break is driven by sector-wide technology, content-cost amortization mechanics, or industry S-curve consolidation, the Netflix series should exhibit a comparable 2018–2019 acceleration. Empirically, Netflix's $\widehat{\gamma}_{p,t}$ is well approximated by a single linear trend with $R^2 = 0.84$; the post–pre level shift of 5.2 percentage points is small relative to the iQIYI shift of 12.1 percentage points. Although the Netflix Chow test attains $p = 0.019$ owing to the low residual variance around a tight linear trend, the absolute magnitude of the Netflix shift is an order of magnitude smaller than iQIYI's, and the Netflix series exhibits no change in slope at the 2018–2019 boundary.