Competition in a Declining Industry: Evidence from Traditional Pay TV

James Atkins

8 November 2026

Click here for the latest version

Abstract

Traditional pay television subscriptions in the United States have nearly halved since 2015, while programming costs have risen sharply and streaming has emerged as a close substitute. I decompose this decline into the effects of rising prices and an improving outside option. Using an event study of broadcast television station acquisitions, I find that TV station acquisitions increase downstream prices by approximately 3.6 per cent over two years, consistent with higher retransmission fees being passed through as itemised charges. I then estimate a random-coefficients demand model with a time-varying outside option. The estimates imply that the outside option improved by the equivalent of a $40 reduction in monthly pay television prices between 2015 and 2025. This quality improvement explains approximately 90 per cent of the decline in cable subscriptions.

James Atkins: Northwestern University, JamesAtkins@u.northwestern.edu.

1 Introduction

Traditional pay television is an industry in decline. After peaking in 2012, traditional pay TV distributors — cable providers such as Comcast or satellite providers such as DirecTV — lost more than half of their US subscribers between 2015 and 2025. Over the same period, streaming services such as Netflix and Hulu became an increasingly attractive alternative while pay TV providers faced sharply rising programming costs. How much of the decline of pay TV reflects higher subscription prices, and how much reflects a shift in the relative attractiveness of traditional television and its outside option?

I find that the increasing relative attractiveness of the outside option, not prices, drives the decline. Relative to traditional pay television, the value of the outside option rises by the equivalent of almost $40 a month between 2015 and 2025. Had it not improved after 2015, 70 per cent of households would subscribe to pay television in 2025 rather than 32 per cent, close to the 80 per cent who subscribed in 2015. The improvement has also changed whom cable competes with. In 2015 satellite was the largest destination of the subscribers a cable system lost to a price increase; by 2025, 76 per cent of them leave pay television altogether, up from 27 per cent.

These findings bear on the current debate over consolidation among broadcast station groups. The Federal Communications Commission (FCC) has recently relaxed its national ownership cap, and the merger of Nexstar and TEGNA, the largest and third-largest owners of television stations, has received federal approval. Most Big Four (ABC, CBS, NBC and Fox) affiliates belong to station groups that own stations across many local markets and negotiate retransmission fees for their whole portfolio with each distributor. An acquisition brings the target's stations under the buyer's agreements and may strengthen the buyer's bargaining position. The case that consolidation harms households rests on the resulting fees reaching subscribers' bills, and on how much of them distributors can pass on given the demand they face. I provide evidence on both. An event study of station acquisitions shows that upstream ownership changes reach downstream receipts per subscriber in the acquired station's market. The demand estimates show that the households over which station groups bargain increasingly choose an outside option that includes services retransmission consent does not govern.

I estimate a random-coefficients discrete-choice model in which households choose among locally available cable distributors, satellite television and an outside option that includes streaming and over-the-air television. The value of the outside option relative to traditional television changes over time through a national index common to all households, and varies across households with their demographics and local broadband availability. I combine subscriber totals with survey moments from the NTIA Internet

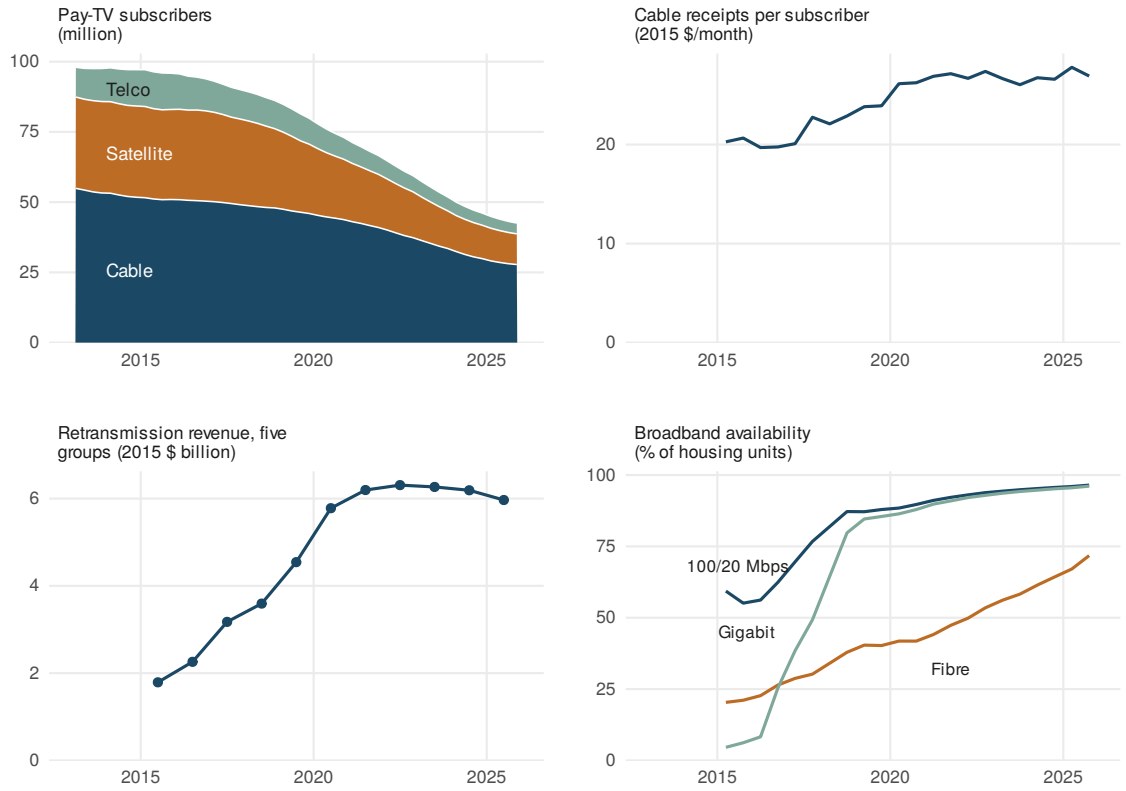

FIGURE 1. Pay-television subscribers, cable receipts per subscriber, retransmission revenue and broadband availability. Subscribers are Kagan's national counts of cable, satellite and telco video subscribers. Receipts per subscriber is the Section 111 gross receipts of the major MVPDs' cable systems divided by their Kagan subscribers. Gross receipts are the charge for the broadcast-basic tier, which every subscriber buys, including broadcast fees, together with additional-set and converter fees; they exclude higher tiers, premium channels, installation, late fees and non-video services. Retransmission revenue is the total reported in the 10-K filings of Nexstar, Gray, TEGNA, Scripps and Sinclair. Broadband availability is the share of housing units in census blocks where the FCC records a provider offering 100/20 Mbps, gigabit (940 Mbps or more) or fibre service. Dollar amounts are in 2015 dollars.

Use Survey and MRI-Simmons on the demographic composition of subscribers. The national index is the common time path of cable's mean utilities once prices are accounted for. It therefore measures the outside option relative to traditional television: it does not separate improvements in streaming from a decline in the appeal of pay television common to all distributors.

Consistently defined subscription prices are not observed across distributors and markets, so I measure the price of basic service by gross receipts per subscriber, which each cable system reports twice a year to the Copyright Office under Section 111 of the Copyright Act. Receipts count only tiers that carry a broadcast signal, and federal law confines broadcast signals to the basic tier that every subscriber must buy, so receipts

per subscriber are approximately the basic charge: the basic tier's rate-card price plus any broadcast fee or surcharge. Every subscriber pays it, so it is the price of the decision to subscribe. The subscriber counts in the denominator are measured with error, so I instrument receipts per subscriber with the same receipts divided by an independently constructed subscriber count. This instrument corrects the measurement error but not the response of the basic charge to local demand, which biases price sensitivity towards zero. A shift-share instrument that interacts national retransmission fee paths with predetermined local exposure to station groups supplies an overidentifying restriction, but has little first-stage power once DMA effects are absorbed. Standard errors are clustered by system and account for the estimation of the household heterogeneity parameters.

Estimating demand raises a computational problem: my subscriber data is at different levels of aggregation. Cable systems have irregular, often overlapping, footprints while satellite subscribers are reported for broader geographic areas. While this could be written as a Berry et al. (1995) inversion — with geographical variation in choice sets of cable systems represented through household heterogeneity — this is computationally inefficient. I instead impose the observed totals as constraints in a mathematical program with equilibrium constraints, following Dubé et al. (2012), and group households into cells that share the same set of available systems. This grouping is exact, and because each cell contains only a few alternatives the programme is sparse and tractable to optimizers.

To separate the contributions of prices and the outside option, I evaluate demand at each combination of their 2015 and 2025 values, holding the remaining features of the 2025 market fixed. Together they imply a fall of 41 percentage points in penetration, against an observed fall of 48 points, and a Shapley decomposition attributes approximately 90 per cent of these 41 points to the improving outside option and 10 per cent to rising prices. Resetting the outside option alone undoes 78 per cent of the observed fall; the two figures share a numerator and differ in their denominators. The improving outside option also changes the competitive constraints facing cable distributors. Holding 2025 prices and choice sets fixed, returning the outside option to its 2015 level reduces the share of a cable system's marginal subscribers who leave pay television from 76 to 45 per cent and lowers the magnitude of the subscriber-weighted own-price elasticity from 1.25 to 0.94. The outside option thus both draws households away from traditional television and makes the remaining demand for cable more price sensitive. These changes bear on distributors' incentives to pass through programming costs, although pass-through also depends on the curvature of demand.

The acquisition analysis exploits changes in the ownership of Big Four affiliates. When a station changes hands it may enter the buyer's carriage agreements, either

through after-acquired-station provisions or at renewal, and the buyer may gain bargaining leverage. I estimate a stacked event study comparing receipts per subscriber in acquired markets with that in other markets served by the same distributor, which absorbs national programming costs and other shocks common to the distributor's operations. Receipts per subscriber rises downstream by approximately 3.6 per cent in acquired markets over the two years after closing. This pattern is consistent with distributors passing higher fees for acquired stations through to subscribers, but the fees themselves are not observed, and the posted basic rate does not respond, so the increase most plausibly runs through the broadcast fee.

This paper contributes to three literatures. The first studies demand, competition and welfare in multichannel television, including Goolsbee and Petrin (2004), Crawford and Yurukoglu (2012), Shcherbakov (2016), Crawford et al. (2018) and Crawford et al. (2019). I model the value of the outside option relative to pay television as changing over time and separate the contribution of its improvement to declining subscriptions from its effect on substitution at the margin. The second is the estimation of differentiated-products demand from aggregate and micro data, following Petrin (2002), Berry et al. (2004) and Conlon and Gortmaker (2026). I show that demand can be estimated when reported quantities aggregate over overlapping choice sets rather than conventional markets, and that a sparse MPEC formulation, including the ratio-type micro moments, makes this tractable at national scale. The third studies bargaining between content owners and distributors and the effects of consolidation, including Chipty and Snyder (1999), Crawford et al. (2018) and Stahl (2016), and in particular cross-market mergers, such as Vistnes and Sarafidis (2013) and Dafny et al. (2019). Comparing acquired markets with those where the buyer already owned stations, I document the geographic scope of the downstream response to upstream acquisitions without observing proprietary carriage contracts.

Section 2 describes the industry and Section 3 the data. Section 4 presents the event study. Section 5 develops the demand model, Section 6 discusses identification and estimation and Section 7 reports the estimates. Section 8 presents the counterfactuals and Section 9 concludes.

2 Industry Background

2.1 Networks, Stations and Distributors

Broadcast television in the United States has a three-level vertical structure, illustrated in Figure 2. At the top are networks, the most important of which are the "Big Four":

ABC (owned by Disney), CBS (Paramount), NBC (Comcast) and Fox (Fox Corporation). Networks produce or acquire national programming such as prime-time entertainment, national news and major sports events. The networks do not reach households directly. Instead, each network distributes its programming through one local TV station in each of the 210 designated market areas (DMAs) into which Nielsen divides the country. In the largest markets these stations are owned and operated by the network itself; elsewhere they are affiliates, owned by station groups such as Nexstar, Sinclair, Gray, TEGNA and Scripps, which hold long-term affiliation agreements with the network. Each local TV station broadcasts the national network feed alongside its own local news and syndicated programming and sells local advertising.

Households can receive broadcast stations free over the air with an antenna, but most have historically received them through a multichannel video programming distributor (MVPD). MVPDs comprise cable operators, such as Comcast, Charter, Cox and Altice, and the two direct-broadcast satellite providers, DirecTV and DISH. A cable operator serves households through local systems whose footprints follow municipal franchise areas; a satellite provider serves the whole country and has delivered local stations by satellite in almost every DMA since the Satellite Home Viewer Improvement Act of 1999. Both bundle local broadcast stations with cable networks such as ESPN and CNN and sell the bundle as a set of tiers. For cable operators the 1992 Cable Act requires that every subscriber purchase a basic tier containing the local broadcast stations before buying any other programming, so broadcast stations are part of every cable bill.

Two FCC ownership rules shape how station groups have grown. The national ownership cap — overturned in August 2026 — limited the stations a single group may own to those reaching 39 per cent of television households throughout my sample period, and the local ownership rule generally prohibits a group from owning two of the four highest-rated stations in the same DMA[1]. Consolidation has therefore taken place mostly *across* markets: a group expands by acquiring Big Four affiliates in DMAs where it does not already own one, so that a typical transaction changes the owner of a market's station without changing the number of independent owners in that market (Stahl 2016). This is the source of variation exploited in the event study.

[1]The FCC allowed broadcast TV companies to count Ultra-High Frequency (UHF) station audiences at only 50% of their actual household reach when calculating the nationwide 39% ownership cap, so station groups can reach a much larger share of the national audience than the nominal 39% cap implies. Under the Top-Four Prohibition — overturned in July 2025 — a single entity was generally barred from directly or indirectly owning, operating, or controlling two of the top four highest-rated television stations in the same Designated Market Area.

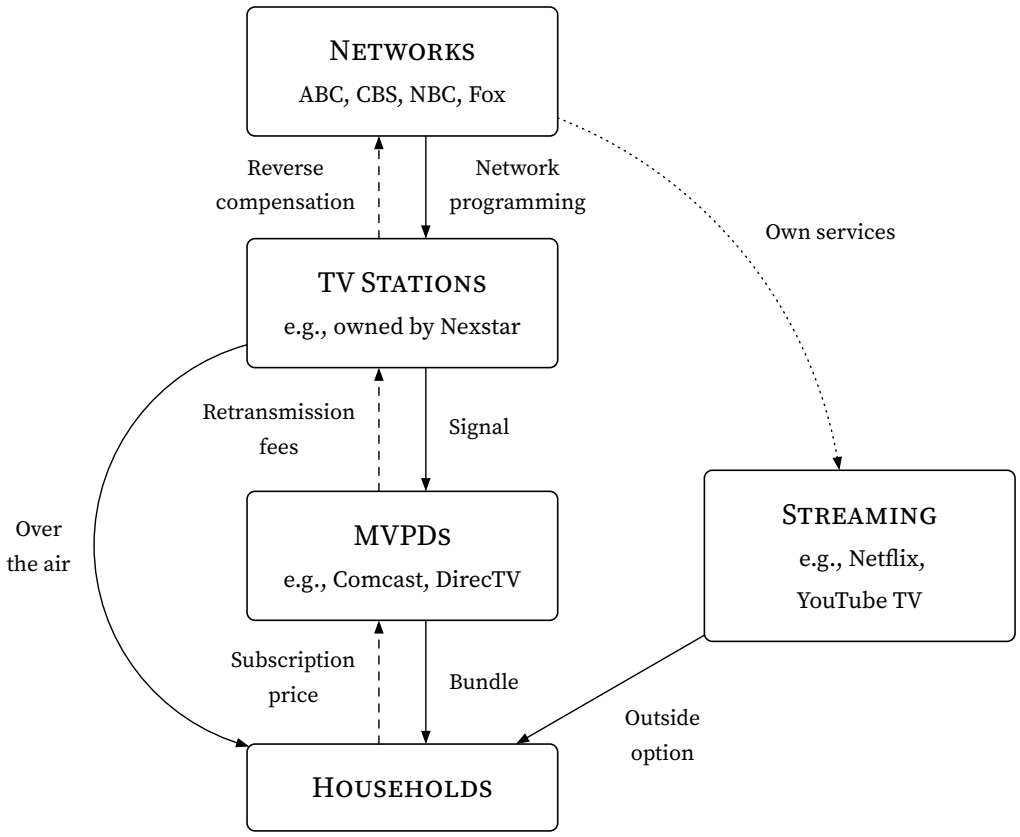

FIGURE 2. Vertical structure of traditional pay television. Solid arrows denote the flow of programming and dashed arrows the flow of payments. The dotted arrow indicates that the networks' parent companies own several streaming services.

2.2 Retransmission Consent

The 1992 Cable Act gave commercial stations a property right in their signals. Every three years each station elects either *must-carry*, under which local MVPDs are obliged to carry it but pay nothing, or *retransmission consent*, under which an MVPD may carry it only with the station's permission. Big Four affiliates almost all elect retransmission consent and negotiate the terms of carriage bilaterally with each MVPD. If negotiations fail, the MVPD loses the right to carry the station and its subscribers experience a *blackout* until an agreement is reached. Until the mid-2000s consent was typically granted in exchange for in-kind compensation, most often the carriage of cable networks affiliated with the station's owner. Since then stations have demanded cash payments set per subscriber per month and these fees, together with advertising, have become a principal source of revenue for station groups. Industry estimates put the average retransmission fee at $5.37 per subscriber-month in 2015 and $21.71 in 2024. Not all of this revenue stays with the stations: through their affiliation agreements the networks claim a share of their affiliates' retransmission revenue, known as *reverse compensation*.

A station group negotiates a single agreement with each MVPD that covers all of the group's stations in the MVPD's footprint. Agreements typically last three years, often aligned with the election cycle, and their terms are confidential. However the industry standard is that these agreements consist of per-subscriber per-month fees for any Big Four stations, and a separate fee schedule for stations affiliated with other networks. Per-subscriber fees therefore vary across station group and MVPD pairs, with larger groups generally commanding higher fees, but no public source reports them at that level. Fees are observed only in aggregate, from station groups' financial filings and industry estimates. This is why the event study measures the consequences of acquisitions for what distributors charge households rather than for the fees that distributors pay.

Cable operators price at the level of an operating region rather than at the household level. For the largest operators the posted rate card is national or follows a legacy footprint (Charter's former Time Warner Cable and Bright House territories, Altice's Optimum and Suddenlink territories), while mid-sized operators price by state or DMA. Satellite providers set national prices. In addition to the posted basic rate, most operators since the mid-2010s itemise a broadcast TV surcharge that recovers retransmission costs, such as Spectrum's "Broadcast TV Surcharge". This surcharge appears in gross receipts but not always in posted basic rates, which bears on the choice of price measure below. Local demand shocks therefore cannot explain price variation across a distributor's systems in the same pricing region, and price responses to local cost shocks arrive through the regional rate card.

2.3 Station Acquisitions and After-Acquired Clauses

As retransmission consent fees are negotiated between station groups and MVPDs, a change in a station's owner can change the fee that each MVPD pays to carry it. An acquired station leaves the seller's agreements and must eventually be brought under the acquirer's. When this happens depends on the contracts in place. Retransmission agreements commonly contain *after-acquired station* clauses, which specify the terms that apply to stations the group acquires during the term of the agreement. Where the acquirer already has an agreement with the MVPD containing such a clause, the acquired station can move to the acquirer's rates as soon as the transaction closes. Where it does not, the seller's agreement typically remains in force until it expires and the acquirer's terms apply only at renewal. If the buyer's applicable fee exceeds the seller's, the transition raises an MVPD's programming costs in the affected DMA either immediately or when the existing agreement expires. The timing can differ across MVPDs serving the same market.

Acquirers anticipate these increases. In its acquisition of Tribune Media in 2019, for example, Nexstar projected annual net retransmission synergies of approximately $75 million, reflecting the application of its own rates to Tribune's stations. Whether and when a given acquisition raises fees is therefore an empirical question; since fees are not observed, Section 4 examines its consequence for receipts per subscriber.

2.4 Streaming

Over the sample period the alternatives to traditional pay television changed considerably. In 2015 the main subscription video-on-demand (SVOD) services were Netflix, Hulu and Amazon Prime Video, which offered libraries of films and television series but little live programming. Between 2019 and 2021 the owners of the major networks and cable networks launched or expanded services of their own: Disney+ (November 2019), HBO Max (May 2020), Peacock (July 2020) and Paramount+ (March 2021, replacing CBS All Access), alongside Apple TV+ (November 2019). These firms increasingly reserved new programming for their own services, and much of the networks' prime-time programming became available on them the day after broadcast.

A second group of services replicates the pay-television bundle itself. Virtual MVPDs, such as Sling TV (launched in 2015), YouTube TV and Hulu + Live TV (both 2017), deliver linear cable networks and, in most markets, the local Big Four affiliates over broadband. They are not MVPDs under the FCC's rules, so the retransmission consent regime does not govern their carriage of broadcast stations; in practice the networks typically negotiate this carriage on behalf of their affiliates. Live sports, long the most defensible content in the bundle, also became available over broadband piecemeal, for example through Amazon's rights to Thursday Night Football from 2022 and the move of NFL Sunday Ticket to YouTube in 2023. Virtual MVPDs do not file under Section 111, and I treat them as part of the outside option together with SVOD services and over-the-air reception.

All of these services require a broadband connection, so their value to a household depends on local broadband availability, which varies across census blocks and over time. In the structural model availability shifts the value of the outside option across households; its common change over time is a separate national index (Section 6.1).

3 Data

Copyright Office Section 111 Statements of Account

My data on cable systems come from filings to the U.S. Copyright Office. Section 111 of the Copyright Act of 1976 permits cable systems broadcast TV signals without individual negotiations with every copyright owner, provided they follow strict statutory rules. Instead of negotiating with each copyright owner, a system pays a royalty computed from its gross receipts from subscribers. To determine the royalty, every system files a Statement of Account (SOA) twice a year, and the Copyright Office examines each filing. Wireline entrants such as AT&T U-verse and Verizon Fios file as cable systems. DirecTV and DISH retransmit under Sections 119 and 122 and do not file.

I collect the filings from two sources. I scrape the Copyright Office's Licensing Division System (LDS)[2] database. This yields, for each system and half-year, the legal owner, the communities served, gross receipts and the number of subscribers. For the largest MVPDs I also collect the statements themselves[3] and process the scanned forms using optical character recognition. Examples are reproduced in Online Appendix B.1.

Gross receipts are my measure of price. They are the revenues over the six-month period from every tier of service that carries a broadcast signal, together with additional-set and converter fees, and they include any broadcast fee that an operator bills outside the base rate. They exclude installation and reconnection charges, late fees, separately billed security or alarm services and programming sold outside a tier that carries a broadcast signal.

Which tiers carry a broadcast signal is determined by law and by how operators have responded to it. The operator defines its tiers, but since the 1992 Cable Act (47 U.S.C. 543(b)(7)) the basic service tier must contain every must-carry and retransmission-consent signal, and every subscriber must buy it. Operators therefore made broadcast basic the only tier whose lineup contains a broadcast station. Expanded tiers require basic but are sold as separate tiers that do not themselves carry a broadcast signal, so their charges fall outside gross receipts. The gross receipts of system k are therefore

$$\text{Gross Receipts}_k \approx p_k^{\text{Basic}} \sum_{\ell} q_{k\ell},$$

where $q_{k\ell}$ is the number of subscribers whose highest tier is ℓ, $\sum_{\ell} q_{k\ell}$ is the system's subscribers and p_k^{Basic}, which I call the basic charge, is the rate-card price of the basic tier plus any broadcast fee or surcharge billed alongside it. Receipts per subscriber are thus

[2] See https://licensing.copyright.gov/lds/.

[3] See https://www.copyright.gov/licensing/111/efilings.html and https://www.copyright.gov/licensing/111/post-examined-soa/.

approximately p_k^{Basic}; the approximation omits additional-set and converter fees. I call this ratio receipts per subscriber rather than average revenue per subscriber: it excludes higher tiers, premium channels and broadband, so it measures the basic charge and not what the average subscriber spends. They do not depend on how subscribers divide among higher tiers.

The LDS subscriber count is the number of "first sets", the primary connection in a household as opposed to additional outlets. It is a quantity defined by the royalty calculation, not a count of video subscribers, and its coverage deteriorates over the sample. Between 2015 and 2024 Charter's first-set count falls from 94 to 64 percent of the video subscribers it reports to investors, and Cox's counts are unusable from 2022. I therefore take quantities from Kagan and use the first-set count only as a check.

Space E of the SOA lists the rates a system charges for service to the first set. Systems typically report a range, and I take the lowest rate to be the price of the basic tier. This rate has two shortcomings. Most filers bill the broadcast surcharge as a separate line item that Space E omits. The rate is also subject to filer-specific reporting changes: Charter reports $9.99 for almost every system in every year from 2018 to 2024. I therefore use the Space E rate only as a robustness check.

Each filing lists the communities in which the system retransmits broadcast signals. I match these communities to Census places, county subdivisions and counties using the Census Gazetteer files, which gives the footprint of each system in each half-year. Table 1 summarises the filings.

FCC Broadband and Cable Registration Data

No single source records which MVPD serves which household. The community lists in the SOAs record where an operator holds a franchise; I combine them with the FCC's broadband availability filings, which record where an operator's plant can deliver service, and with the FCC's cable registration system, which records homes passed and transmission technology.

Broadband availability comes from FCC Form 477 through December 2021 and from the Broadband Data Collection (BDC) thereafter. Both report the providers offering fixed broadband in each census block, with the technology and speeds offered. They measure the broadband available to a household, which shifts the value of the outside option, and through the provider identifier they record where each cable operator's plant reaches. The BDC reports availability by location rather than by block, and the change of source appears in a block-level panel as widespread exit in 2022. I therefore do not interpret changes across the two sources as entry or exit.

TABLE 1. Cable systems filing Section 111 Statements of Account. Each column is a first half-year. Dollar amounts are in 2015 dollars. Gross receipts and subscriber counts are totals over systems. Large systems file Form SA3, which the Copyright Office requires above a gross-receipts threshold. Kagan subscribers in panel A are Kagan's national cable and telco total. In panel B they are Kagan's county estimates apportioned to systems. Kagan does not identify the other filers, so panel C is the difference. A system's franchise area is the Census geography matched to the communities it lists; its housing units are Census block counts. The posted basic rate is the lowest Space E rate for service to the first set.

	2015/1	2020/1	2025/1
A. All Section 111 filers			
Systems	3,197	2,686	2,229
Gross receipts ($ billion)	9.69	9.76	6.35
Kagan subscribers (million)	64.3	53.2	32.6
Gross receipts per system, median ($ thousand)	175	223	153
Franchise-area housing units per system, median (thousand)	5.9	8.9	11.6
Large systems (Form SA3), share of systems (%)	30.1	36.2	33.8
Large systems (Form SA3), share of receipts (%)	97.0	97.7	97.8
B. Systems of major MVPDs			
MVPDs	18	15	15
Systems	1,368	1,258	1,188
Gross receipts ($ billion)	8.80	9.12	6.00
Kagan subscribers (million)	58.3	49.0	30.1
Gross receipts per system, median ($ thousand)	454	562	309
Franchise-area housing units per system, median (thousand)	19.4	23.7	24.1
Kagan subscribers per system, median	4,092	3,911	1,834
Systems with a Kagan estimate	1,137	1,045	1,024
Posted basic rate, median ($/month)	28.33	27.42	29.44
Systems reporting a posted rate	1,049	972	934
C. Other filers			
Systems	1,829	1,428	1,041
Gross receipts ($ billion)	0.89	0.64	0.36
Kagan subscribers (million)	6.0	4.2	2.5
Gross receipts per system, median ($ thousand)	113	130	85
Franchise-area housing units per system, median (thousand)	2.4	4.0	5.2

The Cable Operations and Licensing System (COALS) holds two forms that I use. On Form 322 an operator registers each community it serves, together with the community's county and the operator's mailing address. The county distinguishes communities that share a name, and the address links filers reporting under a subsidiary's name to their parent. Form 325, the annual cable system report, records homes passed, miles of coax-

ial and fibre plant and channel capacity; the FCC discontinued it after 2016. I do not use COALS to construct footprints. Registrations are rarely updated after a system changes hands, so they cannot date when an operator began or stopped serving a community.

S&P Kagan MediaCensus

S&P Global's Kagan MediaCensus reports quarterly estimates of multichannel video subscribers from 2013 onwards. For each DMA it reports households, television households and over-the-air households, together with cable, telco and direct broadcast satellite subscribers and their split across the largest operators and packages. A county-level file reports the same counts for 941 operators, although it omits DirecTV before 2021. Kagan counts each operator's subscribers consistently across markets and over time, which the first-set counts in the SOAs do not, and it is the only source of satellite subscribers below the national level. Aggregate Kagan subscriber numbers are very close to subscriber numbers that public MVPDs publish in their SEC filings.

Kagan does not report subscribers by cable system, so I back them out. The Section 111 filings are a strong predictor of Kagan's counts: across firm-by-DMA cells whose systems lie within one DMA, a regression of log Kagan subscribers on log first-set subscribers with firm-by-period fixed effects has a slope of 0.89 and an R^2 of 0.91. This is unsurprising, since the SOAs are one of the few public sources of subscriber counts by system. The filings are informative about the relative size of a firm's systems but not about their level or time path, which drift as filers consolidate or split their statements.

I therefore reverse the relationship: I distribute Kagan's county counts across the systems serving each county in proportion to their filed counts. Kagan then sets the level and time path of each operator's subscribers, and the filings only how they divide among its systems. The construction is described in Online Appendix A.1.

Station Ownership and Affiliation

From the Broadcasting Yearbook, currently published by Grey House Publishing, and FCC ownership and transfer-of-control records I assemble a panel of every full-power commercial station's owner and network affiliation from 2013 onwards. Transfers of Big Four affiliates between station groups, with the consummation date of the transaction, define the event-study treatments. The largest transactions in the sample are Nexstar–Media General (2017), Sinclair–Bonten (2017), Gray–Raycom (2019), Nexstar–Tribune (2019), Scripps–Cordillera (2019), Gray–Meredith (2021), and the TEGNA divestitures. Figure 3 illustrates the number of stations that change owner each half-year.

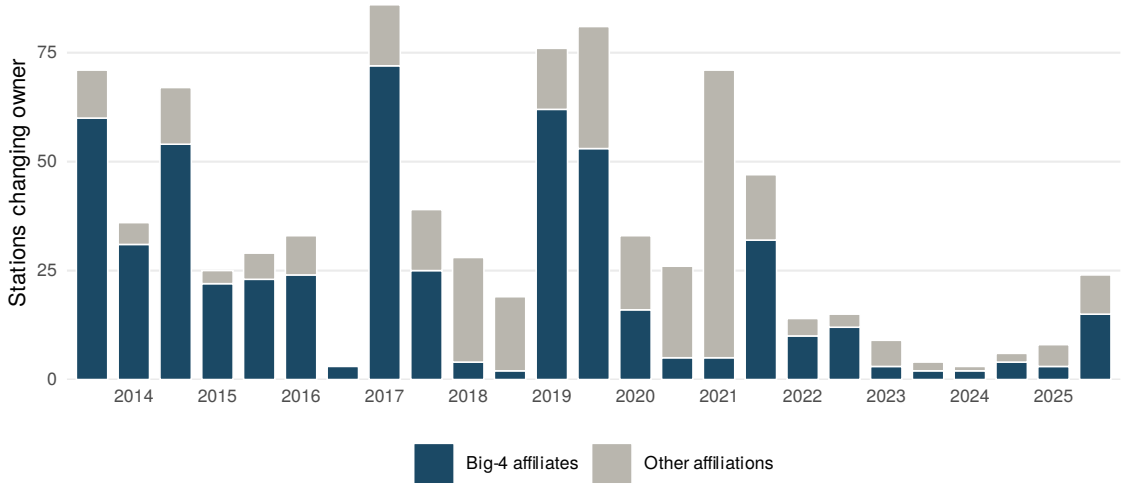

FIGURE 3. Full-power commercial stations changing owner group, by half-year, from 2013/2. A station's affiliation is the one it carried before the transfer. Owner groups attribute a sidecar licensee to the station group that operates it.

NTIA Internet Use Survey

The survey is a supplement to the Current Population Survey fielded in 2015, 2017, 2019, 2021 and 2023. The full microdata report whether the household subscribes to traditional pay TV together with age, income, rental status and broadband status. The baseline uses moments on the demographic composition of pay-television subscribers.

MRI-Simmons

Annual cross-tabulations for 2015–2024 report joint and conditional probabilities of cable, satellite and no-pay-TV status by income, age, rental status, census division and Nielsen county size. Combined with DISH's SEC filings they also yield regional satellite subscriber counts. The baseline matches Kagan's DMA-level satellite totals and the NTIA demographic moments; the MRI-Simmons composition of satellite subscribers by county size enters a robustness specification (Section 6.1).

Household demographics. Households are simulated at the census block-group level. The ACS reports pairwise margins of age, income, and rental status at the block-group level; PUMS reports the full joint distribution at the PUMA level. Entropy balancing finds the block-group joint distribution closest to the PUMA joint distribution subject to the ACS margins, which yields integration weights over a common set of demographic nodes (Appendix B).

4 Station Acquisitions and Receipts per Subscriber

A station acquisition changes the counterparty from which a distributor buys the right to carry a local broadcast signal. It can raise the fees a distributor pays in two ways. The acquired station may move onto the buyer's existing carriage terms, either at closing under an after-acquired clause or when the seller's agreement expires. The acquisition may also strengthen the buyer's position in later negotiations, raising fees for its other stations as well as the acquired one. Neither a higher fee nor its transmission to subscriber bills follows from a change in ownership alone: the first depends on the contracts and the bargaining between station groups and distributors, the second on downstream demand and pricing.

I examine the downstream consequences of upstream consolidation using acquisitions of Big Four affiliates. I find receipts per subscriber rises by about 3.6 per cent in acquired markets relative to other markets served by the same distributor over the two years after closing. This estimate is robust across across specifications that vary the definition of an acquisition, distinguish the buyer's existing markets and allow for a response common to an exposed distributor's markets. In this section, I first establish the baseline result, then examine its robustness to these changes in the sample and comparison group.

4.1 Research Design

I use the history of station acquisitions, as shown in Figure 3, as the basis of my event study. I define an acquisition as a change in the ultimate ownership of a Big Four affiliate in television market m, dated by the closing of the transaction as reported by the FCC. Let d index a distributor and t a half-year, which I call a period. For each closing half-year c, I compare distributor–market pairs in which at least one affiliate changes hands with pairs in markets without an ownership change in the four half-years before or after c. A treated market is excluded if another transfer occurs in a different half-year of this window. Several affiliates can change owners in the closing half-year, so the treatment is a broader "acquisition episode" rather than necessarily a single station or transaction.

Let R_{dmt} denote half-year gross receipts and N_{dmt} Kagan's subscriber count. The outcome variable is log monthly receipts per subscriber, $y_{dmt} = \log(R_{dmt}/6N_{dmt})$. Receipts reported under Section 111 cover the charge for the broadcast-basic tier that every subscriber buys, together with equipment charges and itemised broadcast surcharges, but not the charges for higher tiers (Section 3). Where a cable system spans television markets, I allocate its receipts in proportion to estimated subscribers (Appendix Equation 16).

I stack the cohort samples and estimate the event study

$$y_{dmt} = \mu_{dmc} + \lambda_{dct} + \sum_{k=-4, k\neq -1}^{4} \theta_k \mathbb{1}[t - c = k] D_{mc} + \varepsilon_{dmtc}, \tag{1}$$

where D_{mc} equals one if an affiliate in market m changes owners in cohort c. The fixed effect μ_{dmc} allows each distributor–market pair a separate level in each cohort, while λ_{dct} absorbs changes common to a distributor's markets within that cohort. The same observed outcome can enter several cohort samples. Each θ_k measures the change in the acquired–comparison market difference relative to the half-year before closing, $k = -1$. I cluster standard errors by television market to allow dependence over time and across repeated appearances in the stack. The sample contains 115 treated and 181 comparison markets over 17 cohorts (Appendix Table 10).

The distributor-by-time effects are central to the design. They remove movements in log receipts per subscriber common to a distributor's markets, including common national pricing and programming-cost shocks. Identification therefore comes from by comparing the differences in receipts per subscriber between markets served by the same distributor at the same date. The identifying restriction is that, absent the acquisition, these within-distributor differences would have followed parallel paths. A uniform distributor-wide response is absorbed by λ_{dct}, so the coefficients measure the local component of the response rather than the acquisition's total effect on the distributor's customers. Unrelated local shocks correlated with acquisitions would violate this identifying restriction.

Transactions are announced before they close, a median of 10 months earlier, and 49.2% of announcements fall in $k = -2$. The reference period $k = -1$ may therefore already contain anticipatory changes. The alternative normalisation $k = -3$, a robustness check for these anticipation effects, is reported in Appendix Table 11.

4.2 Results

Figure 4 reports the estimates of Equation 1. The pre-closing coefficients are jointly insignificant ($F = 0.47$, $p = 0.70$) and there is no increase in the closing half-year. The response then rises from 0.033 log points at $k = 1$ to 0.070 at $k = 4$. This timing is consistent with staggered contract expiry or delayed retail price adjustment. It does not distinguish them: individual renewal dates are unobserved, and an immediate wholesale fee change need not produce an immediate retail response.

To summarise the post-closing response, I replace the event-time interactions in Equation 1 with a single interaction $D_{mc}\mathbb{1}[1 \leq t - c \leq 4]$ and omit the closing half-year. Table 2 reports this pooled difference-in-differences coefficient, comparing the

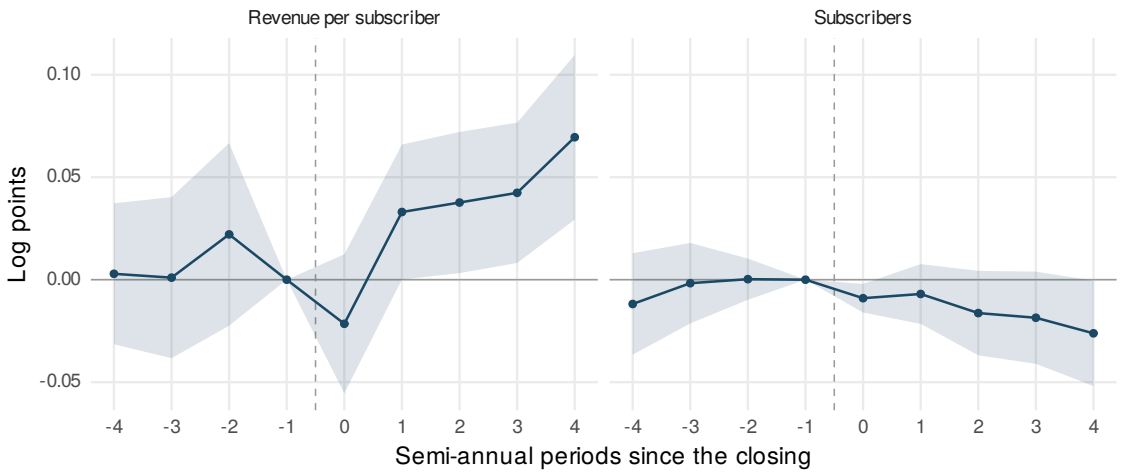

FIGURE 4. Effects of Big Four station acquisitions on log receipts per subscriber (left) and log subscribers (right). Coefficients θ_k from Equation 1 relative to the half-year before closing, with distributor × cohort × period effects. Bands are 95 per cent confidence intervals from standard errors clustered by television market.

four post-closing half-years with the four pre-closing half-years. It is not an arithmetic average of the plotted coefficients, which use $k = -1$ as their reference. The baseline estimate in column (1) is 0.036 (0.014). Column (2) replaces the indicator with the share of the market's Big Four affiliates acquired; at the mean treated share of 0.35 it implies 0.026, against 0.036 for the indicator. Column (3) replaces the distributor-by-time effects with cohort-by-time effects and gives a slightly larger 0.041.

TABLE 2. Pooled post-acquisition coefficients, in log points, comparing $k = 1, \ldots, 4$ with $k = -4, \ldots, -1$ and excluding $k = 0$. All specifications include distributor–market × cohort effects; distributor-time effects are distributor × cohort × period, and columns without them include cohort × period effects. The share coefficient in column (2) is for a change from no Big Four affiliates acquired to all of them. Column (4) minus column (5) equals column (1). Column (6) is the log posted basic rate, the lowest monthly first-set rate in Space E of the Statements of Account; where several systems serve a distributor–market it is their average weighted by footprint housing units, over the systems that report one. Standard errors clustered by television market are in parentheses. Alternative post-periods are reported in Appendix Table 12.

	Log receipts per sub.			Log receipts	Log subs.	Log basic rate
	(1)	(2)	(3)	(4)	(5)	(6)
Acquisition	0.036		0.041	0.023	−0.014	−0.011
	(0.014)		(0.014)	(0.020)	(0.013)	(0.013)
Share of affiliates acquired		0.075				
		(0.031)				
Distributor-time effects	Yes	Yes	No	Yes	Yes	Yes
Observations	27,305	27,305	27,312	27,305	27,305	22,559

Columns (4) and (5) separate receipts from subscribers. Because the outcome is $\log R - \log N - \log 6$ and both regressions use the same sample and regressors, the receipts coefficient minus the subscriber coefficient equals column (1). The point estimates are 0.023 for receipts and −0.014 for subscribers. Thus both higher receipts and a lower subscriber count contribute to the estimated increase in the ratio — as expected by a downwards sloping demand curve — although neither component is individually precise. The advertised entry-tier rate in column (6) shows no detectable response. This is consistent with an adjustment through broadcast surcharges or other payments omitted from the posted rate.

4.3 After-Acquired Clause vs Bargaining

The baseline compares acquired markets with all other markets served by the same distributor, including those in which the buyer already owned a station. I now separate these markets to ask where the increase occurs. The answer bears on how an acquisition changes carriage terms. An acquired station that moves onto the buyer's existing rates raises a distributor's costs in that station's market alone, whereas a stronger bargaining position for the buyer can raise them wherever the distributor carries the buyer's stations.

I first compare acquired markets with the buyer's existing markets within a distributor. These comparisons see only the local part of any response, because any adjustment common to a distributor's markets is absorbed by the time effects. I then compare distributors to ask whether the increase also reaches the rest of a distributor's footprint, which 15 distributors allow me to do only imprecisely. In both comparisons the increase specific to acquired markets stays close to the baseline estimate.

Specification and Predictions

Because the buyer's other markets are defined relative to a particular buyer, the cohorts here are transactions denoted by e, each with a buyer b_e and a closing half-year c_e; there are 51 in total. A transaction's sample contains the markets in which no Big Four affiliate changes ultimate owner in $[c_e - 4, c_e + 4]$ other than through the buyer's purchases at closing. Markets with any other ownership change in the window, including a further purchase by the same buyer, are excluded. Within this sample I classify markets as acquired, portfolio or outside markets,

$$\begin{aligned} D^{A}_{me} &= \mathbb{1}[b_e \text{ acquires a Big Four affiliate in } m \text{ at } c_e] \\ D^{P}_{me} &= \left(1 - D^{A}_{me}\right) \cdot \mathbb{1}[b_e \text{ owns a Big Four affiliate in } m \text{ at } c_e - 1] \\ D^{O}_{me} &= 1 - D^{A}_{me} - D^{P}_{me}. \end{aligned} \tag{2}$$

In an acquired market the buyer purchases at least one Big Four affiliate at closing. In a portfolio market it already owned one before closing and no station changes hands. In an outside market it owns none before closing and acquires none: these markets are the comparison group. A transaction has on average 2.7 acquired and 17.8 portfolio markets, since most buyers are established station groups.

Whether the acquisition changes a distributor's dealings with the buyer depends on the distributor's footprint. Let

$$E_{de} = \mathbb{1}[d \text{ serves a market in which } b_e \text{ acquires a Big Four affiliate}] \tag{3}$$

denote whether distributor d is exposed to the transaction, with its footprint measured at $c_e - 1$. For every distributor–market (d, m) in the sample, $D^A_{me} = 1$ implies $E_{de} = 1$, since d then serves an acquired market. Portfolio and outside markets can belong to either kind of distributor. The sample contains 899 exposed and 772 unexposed portfolio distributor–market–transactions.

I then estimate the pooled specification

$$y_{dmt} = \mu_{dme} + \lambda_{det} + \mathbb{1}[t > c_e]\{\beta^A D^A_{me} + (\beta^P_1 E_{de} + \beta^P_0 (1 - E_{de}))D^P_{me}\} + \varepsilon_{dmte} \tag{4}$$

over the four half-years on either side of closing, omitting $t = c_e$. The effect μ_{dme} gives each distributor–market its own level in each transaction's sample, and λ_{det} is a distributor × transaction × period effect. Each coefficient therefore compares a market with the same distributor's outside markets in the same half-year. I also estimate an event-time version, replacing $\mathbb{1}[t > c_e]$ with an indicator for each half-year relative to $k = -1$.

Under *rate harmonisation* the acquired station moves onto the buyer's existing terms, either through an after-acquired clause or when the seller's agreement expires. Under *bargaining leverage* the acquisition improves the terms the buyer can negotiate for the acquired station and its existing portfolio. Table 3 sets out what each channel implies for the coefficients. Both raise the acquired station's fee, so both predict $\beta^A > 0$. They differ in the buyer's existing markets. Harmonisation changes only the acquired station's terms and leaves $\beta^P_1 = 0$. Leverage raises the terms the buyer negotiates for the stations an exposed distributor already carries, so $\beta^P_1 > 0$ once those agreements renew. Neither channel changes the dealings of an unexposed distributor with the buyer, so both predict $\beta^P_0 = 0$ and its portfolio markets serve as a placebo.

Acquired and Portfolio Markets

The estimates in Table 3 follow the pattern that harmonisation predicts. Receipts per subscriber rises where the buyer acquires a station and does not rise where it already owned one, whether or not the distributor also serves an acquired market. The estimate

TABLE 3. Predicted and estimated coefficients of Equation 4. The predicted signs assume that distributors pass fee changes through to receipts per subscriber in the market of the station whose fee changed. The estimates compare $k = 1, ..., 4$ with $k = -4, ..., -1$, omitting $k = 0$, on 75,772 observations; each compares a market with the same distributor's outside markets. Standard errors clustered by television market are in parentheses.

Coefficient	Harmonisation	Leverage	Estimate
β^A: acquired markets	> 0	> 0	0.033
			(0.014)
β_1^P: portfolio markets, exposed distributors	0	> 0 at renewal	−0.003
			(0.011)
β_0^P: portfolio markets, unexposed distributors	0	0	0.004
			(0.014)

in acquired markets is close to the baseline of Table 2, so organising the sample by transaction rather than by closing half-year does not change the response.

The portfolio markets of exposed distributors are where the two channels disagree. If the acquisition strengthened the buyer's hand in its negotiations with a distributor, the stations it already owned would share in the gain, and the distributor's receipts per subscriber in those markets would rise with it. The data rule out a response there of more than about half the acquired-market response: the upper limit of the 95 per cent confidence interval for β_1^P is 0.017. The acquired-market response exceeds the portfolio response by 0.036 (0.016), and the difference is more precisely estimated when I cluster by transaction or by both market and transaction (standard errors 0.010 and 0.013).

Leverage predicts higher fees at the buyer's existing stations only when their agreements are renegotiated. Agreements typically last three years, so on average across closing dates about two-thirds of them renew within the two years after closing. Figure 5 shows that the acquired-market response builds over these two years while the portfolio response shows no rise, including in the later half-years when more of the buyer's agreements have renewed. The gradual rise in acquired markets is itself consistent with acquired stations moving onto the buyer's terms as the seller's agreements expire, or with distributors revising retail charges with a lag; the revenue paths do not separate the two.

Two qualifications apply. First, the bound is per market. Exposed portfolio markets hold almost four times as many subscribers as acquired markets, so a portfolio response below the bound could still raise distributors' total receipts by more than the acquired-market response (Online Appendix C.5). Second, the comparison does not exclude leverage confined to the acquired station. A buyer that negotiates new terms for the

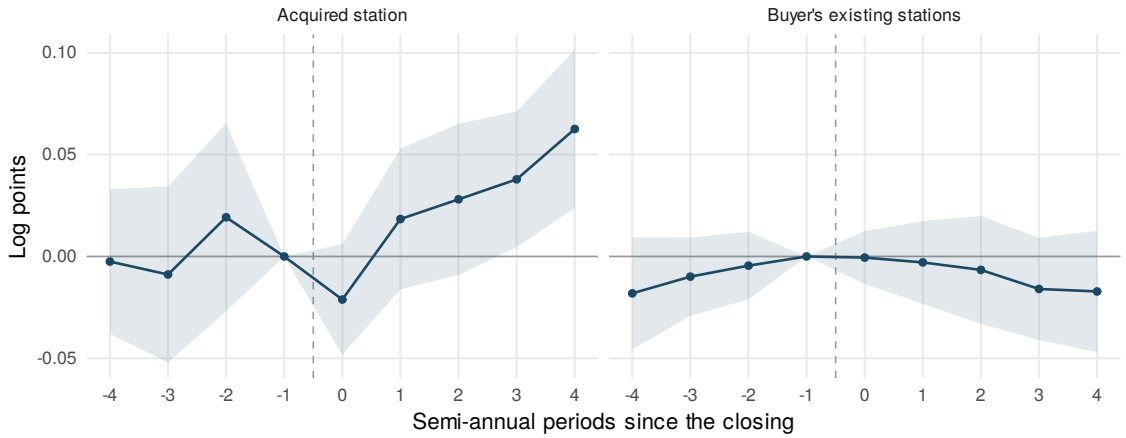

FIGURE 5. Event-time responses in acquired and portfolio markets. Estimates use Equation 4 with distributor × transaction × period effects and reference period $k = -1$. The left panel reports acquired markets and the right panel reports portfolio markets of exposed distributors, each relative to the same distributor's outside markets. Bands are 95 per cent confidence intervals from standard errors clustered by television market.

station it acquires, rather than extending its existing ones, would raise receipts per subscriber in acquired markets alone, just as harmonisation does.

Across Distributors

The largest distributors set one rate card for a region or for their whole footprint (Section 2). A distributor that prices this way spreads a fee increase at any of the buyer's stations over all its markets, and the comparisons within a distributor difference it away. To examine such an increase I replace the distributor × transaction × period effects in Equation 4 with transaction × period effects λ_{et} and add an indicator for the outside markets of exposed distributors:

$$y_{dmt} = \mu_{dme} + \lambda_{et} + \mathbb{1}[t > c_e]\big\{\beta^A D^A_{me} + \big(\beta^P_1 E_{de} + \beta^P_0 (1 - E_{de})\big) D^P_{me} + \beta^O E_{de} D^O_{me}\big\} + \varepsilon_{dmte}. \tag{5}$$

The comparison group is now the outside markets of unexposed distributors, whose dealings with the buyer the transaction leaves unchanged. β^A measures the rise in acquired markets against them, and β^O the rise in the markets of an exposed distributor where the buyer owns nothing, which is the footprint-wide response.

TABLE 4. Acquired-market and outside-market responses. Estimates of Equation 5, comparing $k = 1, ..., 4$ with $k = -4, ..., -1$ and omitting $k = 0$, on 75,803 observations. The first two rows are relative to the outside markets of unexposed distributors; the final row compares acquired markets with the outside markets of exposed distributors. Standard errors clustered by television market are in parentheses.

β^A: acquired markets	0.052
	(0.020)
β^O: outside markets, exposed distributors	0.016
	(0.013)
Acquired minus outside markets, $\beta^A - \beta^O$	0.036
	(0.014)

Allowing a response common to an exposed distributor's markets, the acquired-market estimate is 0.052 (0.020). It comprises a response of 0.016 (0.013) in the distributor's outside markets and a further 0.036 (0.014) specific to acquired markets. The specific part closely matches the baseline estimate of 0.036 in Table 2.

The common component is imprecisely estimated. Exposure varies across only 15 distributors, and clustering by market does not account for shocks common to a distributor's markets. The comparisons across distributors cannot distinguish a common response of zero from one as large as the response specific to acquired markets (Online Appendix C.5). Controls for distributor size leave the acquired-market differential essentially unchanged but alter the estimated outside-market response (Online Appendix C.5).

Additional comparisons using estimated rate differences, relative buyer size and proxies for contractual relationships do not distinguish harmonisation from bargaining gains conclusively (Online Appendix C.6).

I find that upstream station acquisitions raise downstream receipts per subscriber by about 3.6 per cent in acquired markets relative to other markets served by the same distributor over the two years after closing. The acquired-market differential is similar across specifications, while the buyer's existing markets show no comparable differential increase. The evidence thus establishes a local downstream consequence of upstream consolidation, although it leaves open whether distributors also adjust charges across their markets and does not identify the underlying contractual change or its pass-through. Because receipts exclude higher tiers, the increase is not a shift of subscribers towards richer packages: it is a rise in what subscribers pay for basic service, and since the posted basic rate does not respond, most plausibly in the broadcast fee. I next examine the demand conditions governing downstream pricing incentives, focusing on how substitution towards the outside option has changed those conditions.

5 Demand Model

In this section, I model demand for traditional pay television, allowing households to substitute among locally available distributors and an outside option. The model separates two sources of the decline in traditional pay television: movements along the demand curve induced by prices and shifts of the demand curve induced by a rise in the value of the outside option relative to pay television. A better outside option reduces subscriptions even when the prices and quality of cable and satellite service are unchanged, and it alters the substitution patterns that determine how distributors pass programming costs through to subscribers.

The central difficulty in estimating demand is that the alternatives are observed at different levels of aggregation. Cable subscriptions are reported for each system, and Kagan reports satellite subscribers at both the national level and per DMA. A system's footprint need not coincide with a DMA or lie within one. No geographic unit therefore has observed shares for every alternative: the satellite share within a system's footprint and the cable share within a DMA are both unknown. Overlap compounds the problem: cable systems often have overlapping footprints and while I observe subscriber numbers at the system level, I do not observe subscriber numbers in the intersecting portions of the systems' footprints.

Previous work has handled this by discarding data. For example, Crawford et al. (2018) drop systems that overlap with other systems and impute DMA-level satellite shares to the systems that remain. Discarding overlapping systems removes the households for which substitution between cable systems is possible, and imputing satellite shares assumes away differences in satellite take-up between systems in the same DMA.

None of these data need be discarded. I model demand for groups of households that face the same alternatives, which I call choice-set cells, and require predicted subscriptions, summed over the cells in which each alternative is available, to equal its observed total. The allocation of a system's subscribers across the cells it serves, including cells shared with other systems, is then determined by the model rather than imposed. The model therefore maps household choices into the quantities observed in the data without assigning subscribers to overlapping areas in advance. Section 6 describes the identifying variation and estimator.

5.1 Choice Sets

As cable footprints overlap, households served by the same system need not face the same set of alternatives. I therefore compute demand at the level of households that share a common choice set. To do so, I define a *choice-set cell* as a maximal group of

census blocks, fixed over time, such that all blocks lie within the same state and DMA and, in each period, all households face the same set of available cable systems. Every household may in addition choose a satellite provider or the outside option. A cell need not be contiguous, while a cable system may serve several cells and several DMAs. I index cells by g. Formally, choice-set cells are the equivalence class of census blocks under the relation of sharing a state, a DMA and a set of available cable systems in every period, so they partition the census blocks.

Figure 6 illustrates the construction of choice-set cells for two overlapping systems, A and B. Here there are four choice-set cells: $A \setminus B$, $A \cap B$, $B \setminus A$ and $S \setminus (A \cup B)$. Cable prices are common to all households served by the same system, or by the larger pricing region to which it belongs. Households in different cells therefore differ in the systems available to them and in their observed characteristics, but not in the price they face for a given system.

The sample contains 11,987 choice-set cells, fixed over 22 half-years from 2015/1 to 2025/2, spanning 206 DMAs and 49 states. Panel A of Table 5 summarises their size. Cells are small and their size is highly skewed: in 2025/2 the median cell contains 1,466 households and the mean cell 10,702. A cable system serves a median of 3 cells, but the distribution has a long right tail, so the subscription restriction for a large system aggregates over many distinct choice sets.

Panel B classifies cells by the number of cable systems available to them in the final period. In 2025/2, 10.9% of households have no cable system available and choose only between satellite and the outside option. Overlap is the norm among the remainder:

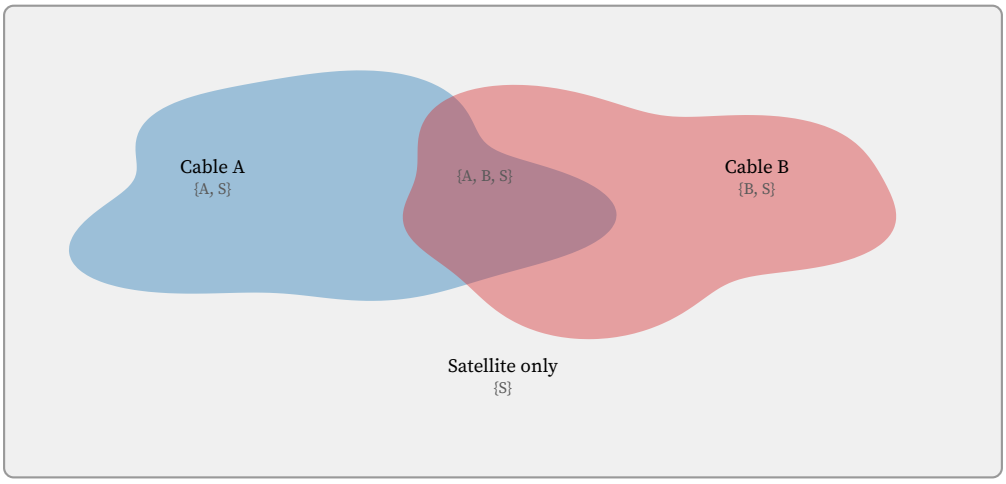

FIGURE 6. Choice-set cells. The footprints of cable systems A and B overlap within a DMA, which partitions its households into four cells with distinct choice sets. Satellite (S) is available in every cell, as is the outside option, which is omitted from the labels. Subscriptions to A and B are observed only as totals over the cells each system serves.

TABLE 5. Choice-set cells. In Panel A, households per cell is measured in the final period, cable systems per cell is computed over cell-periods and cells served per system over system-periods. Panel B classifies cells by the number of cable systems available in the final period. Satellite and the outside option are available in every cell.

Panel A: Size

	10th Percentile	Median	90th Percentile	Mean
Households per cell, 2025/2	151	1,466	19,780	10,702
Cable systems per cell	0	1	2	0.87
Cells served per cable system	1	3	24	10.1

Panel B: Cable systems available, 2025/2

	0	1	2	3	4+
Cells	4,122	5,232	1,993	518	122
Share of households	10.9%	33.3%	37.5%	15.1%	3.2%

61.6% of households with cable can choose between two or more cable operators, and such a household has 1.84 operators available on average. Discarding overlapping systems would therefore remove the majority of cable households from the sample.

5.2 Preferences

In each period t, a household i in cell g chooses between the outside option, the cable systems available in its cell, indexed by k, and satellite[4]. I therefore allow price sensitivity and the value of the outside option, as well as tastes for cable and satellite, to vary across households with income, age and rental status; coefficients indexed by i are random coefficients that can vary this way. Throughout, the taste shocks ε are independently and identically distributed type I extreme value; this yields the familiar logit choice probabilities.

Outside Option

The utility of not subscribing to pay television is

$$u_{igt}^{O} = \lambda_t + \eta_i a_{gt} + \varepsilon_{igt}^{O}, \quad \eta_i = \eta + d_i'\eta^{D}, \tag{6}$$

where d_i contains the household's standardised age, log income and owner-occupation, and a_{gt} is the share of households in cell g with cable or fibre broadband available, measured as a deviation from its mean. The term λ_t is a time-varying index of the value of the outside option, common to all households and normalised to zero in the first half of 2015. A demographic difference in the level of the outside option cannot be separated

[4] I bundle both Dish and DirecTV into a single satellite provider.

from a taste for pay television common to cable and satellite, and enters through β_i^C and β_i^S. The term η_i is the household's value of broadband access. Broadband availability therefore shifts the value of the outside option across cells, and differently for different households through η^D, because streaming services require a broadband connection.

Utility is identified only relative to the inside goods, so λ_t measures the outside option relative to traditional pay television. An improvement in streaming and a decline in the appeal of pay television common to every distributor raise it alike. I refer to it as the value of the outside option in this relative sense, and not as a measure of streaming quality.

Cable

A cable subscriber chooses a programming tier as well as a system. Under the 1992 Cable Act every subscriber must buy the basic tier (Section 3), and higher tiers are sold on top of it, so a household that subscribes to system k buys either basic alone, $\ell = 0$, or basic together with one of the system's higher tiers, $\ell = 1, \ldots, L_k$. Tier ℓ costs $p_{kt}^C + \Delta_{k\ell t}$, where p_{kt}^C is the basic charge (which includes any broadcast fees) and $\Delta_{k0t} = 0$.

I model these choices as a nested logit, with one nest per cable system. The utility of tier ℓ of system k is

$$u_{ik\ell gt}^C = x_{kt}^{C\prime}\beta_i^C - \alpha_i^C\left(p_{kt}^C + \Delta_{k\ell t}\right) + \psi_{k\ell t} + \tilde{\xi}_{kt}^C + \zeta_{ikgt} + (1-\sigma)\varepsilon_{ik\ell gt}, \tag{7}$$

where $\psi_{k\ell t}$ is the value of tier ℓ relative to basic alone, with $\psi_{k0t} = 0$, and the errors follow the nested-logit structure of Berry (1994). The nesting parameter $\sigma \in [0, 1)$ measures the correlation in tastes across tiers of the same system. Integrating over tier choice, a household values system k at the utility of basic alone plus the system's menu value, the value of its higher tiers relative to basic alone:

$$M_{ikt} = (1-\sigma)\log\sum_{\ell=0}^{L_k}\exp\left(\frac{\psi_{k\ell t} - \alpha_i^C\Delta_{k\ell t}}{1-\sigma}\right) = -(1-\sigma)\log s_{\text{Basic}\mid ikt}, \tag{8}$$

where $s_{\text{Basic}\mid ikt}$ is the probability of choosing basic alone, given that the household subscribes to k. Since the basic charge is paid by every subscriber, it shifts the utility of every tier equally and therefore enters system-level demand with coefficient α_i^C. Higher-tier charges and characteristics enter only through the menu value.

I do not estimate tier choice as I only have data on total subscribers per system. The menu value varies across households only through α_i^C: the terms $x_{kt}^{C\prime}\beta_i^C$ are common to every tier of a system and cancel within the nest. With a common price coefficient it is therefore $M_{kt} = -(1-\sigma)\log s_{\text{Basic}\mid kt}$, where $s_{\text{Basic}\mid kt}$ is the share of the system's subscribers who buy basic alone, and I use this as an approximation. The resulting system-level utility is

$$u_{ikgt}^C = x_{kt}^{C\prime}\beta_i^C - \alpha_i^C p_{kt}^C + \xi_{kt}^C + \varepsilon_{ikgt}^C, \quad \xi_{kt}^C = \tilde{\xi}_{kt}^C + M_{kt}, \tag{9}$$

so the system's unobserved quality includes its menu value.

The basic-only share is not observed by system. Kagan reports subscribers by package for each distributor and DMA, and I use that share for each of the operator's systems in the DMA. It varies mostly by operator and period, and only a few operators show substantial variation across DMAs.

The counterfactuals interpret changes in this price measure as changes in the common charge paid across tiers, holding higher-tier increments and characteristics fixed. Changes in the outside option likewise shift the utility of every tier of every system equally. Neither affects the conditional choice among tiers, so the counterfactuals can be evaluated using system-level demand without separately estimating σ.

Satellite

The two satellite providers are combined into a single alternative, with utility

$$u_{igt}^{S} = x_{mt}^{S\prime}\beta_i^{S} - \alpha_i^{S} p_t^{S} + \xi_{mt}^{S} + \varepsilon_{igt}^{S}, \tag{10}$$

where m is the DMA containing cell g. Satellite characteristics and unobserved appeal are defined for each DMA and period, but the price p_t^S is set nationally and so cannot respond to local demand shocks.[5]

5.3 Aggregate Quantities and Mean Utilities

Let $\boldsymbol{\delta} = \left(\boldsymbol{\delta}^{C}, \boldsymbol{\delta}^{S}\right)$ denote the mean utility of cable and satellite respectively, and the θ be the parameters which determine the non-linear component of utility μ. Let $s_{kgt}^{C}(\delta, \theta)$ and $s_{mgt}^{S}(\delta, \theta)$ denote the cable and satellite choice probabilities in cell g, obtained by integrating household choice probabilities over the distribution of household characteristics in the cell. Let H_{gt} be the number of households in cell g, and let $\mathcal{G}_{kt}$ be the set of cells served by system k in period t. Predicted subscriptions must equal the observed totals Q_{kt}^{C} and Q_{mt}^{S}:

$$Q_{kt}^{C} = \sum_{g \in \mathcal{G}_{kt}} H_{gt} s_{kgt}^{C}(\delta, \theta) \tag{11}$$

for every cable system k and period t, and

$$Q_{mt}^{S} = \sum_{g \in \mathcal{G}_{m}} H_{gt} s_{mgt}^{S}(\delta, \theta) \tag{12}$$

for every DMA m and period t, where $\mathcal{G}_m$ is the set of cells in DMA m and Q_{mt}^{S} is the combined subscriber count of the two satellite providers. Equation 11 is the central restriction of the model. Each observed system total constrains the sum of predicted

[5] Satellite providers also sell tiers. The argument for cable applies, and the menu value of satellite's tiers is absorbed by its DMA $\times$ period mean utilities.

subscribers over all the cells the system serves, whereas the allocation of those subscribers across cells, including cells shared with other systems, is not observed and is determined by the model.

The quantity restrictions can be represented as demand in a single national market in each period by giving unavailable cable systems utility $-\infty$ for the households they do not serve. Each observed total is then a market-level quantity and differences in availability are household heterogeneity. The resulting demand system satisfies the connected-substitutes condition of Berry et al. (2013), so for given θ at most one vector of mean utilities δ solves Equation 11 and Equation 12.

6 Identification and Estimation

I estimate the model by combining the aggregate quantity restrictions with survey moments on subscriber composition and instruments for the basic charge. The survey moments discipline household heterogeneity, while the price instruments determine the conversion of utility into dollars. This section first discusses identification of the parameters of interest, and then details my estimation procedure.

6.1 The Outside Option and Household Heterogeneity

I use two sets of micro-moments: one uses MRI-Simmons crosstabs to discipline the cable–satellite margin, the other uses NTIA micro-data to discipline the inside–outside good margin. I use the NTIA Internet Use Survey as micro-moments to match the age, income and tenure composition of households subscribing to traditional pay television. These moments discipline demographic differences in price sensitivity, tastes for pay television and the valuation of broadband-dependent alternatives. The NTIA survey only has a variable for traditional TV and does not distinguish cable from satellite, while the MRI-Simmons cross-tabulations do. I have micro-moments that match the age, income and tenure composition for households subscribing to cable TV only. Satellite has a separate mean utility for each DMA and period, so its geographic totals are matched by Equation 12. In an alternative specification, I impose a nationwide mean utility for satellite and use geographic micro-moments to match the composition of satellite subscribers across the four Nielsen county-size classes, a measure of how urban or rural a county is.

Each micro moment conditions on the choice variable, such as $\mathbb{P}(\text{Income} \mid \text{Pay TV})$, rather than on a household characteristic, such as $\mathbb{P}(\text{Pay TV} \mid \text{Income})$. As recommended by Conlon and Gortmaker (2026), this is to avoid issues with incompatibilities between the survey and aggregate data. The aggregation restrictions already fix the total

market shares; conditioning the survey moments on subscription uses the survey to discipline who subscribes without requiring the model to match a second measure of how many subscribe. The moments are given formally as ratios of survey averages in Online Appendix D.

The national index λ_t is identified instead by the period effects τ_t of the cable equation (Section 6.2), as $\lambda_t = -\tau_t$, under the restriction that the unobserved quality of cable, ξ^C_{kt}, has no common national drift. That restriction is what attributes the common decline of pay television to the outside option rather than to traditional television, and it is why I interpret λ_t as the relative value of the outside option vs pay TV. The survey micro-moments nonetheless condition the estimate, because the nonlinear parameters move the mean utilities that solve the aggregation restrictions.

The broadband terms are identified by cross-sectional variation: how cable and satellite totals, and the composition of subscribers, differ across cells with different availability. This requires availability rather than household broadband adoption, which is itself a choice correlated with demand for pay television. It also requires that the local taste for pay television, conditional on the controls below, does not vary with broadband availability.

6.2 Price Endogeneity and Fixed Effects

The price is gross receipts per subscriber, as in the event study. The unobserved quality ξ^C_{kt} contains local changes in channel lineups, promotional activity, the quality of the operator's broadband service where it is bundled, and the local taste for pay television. It also contains the menu value M_{kt} of the system's higher tiers (Equation 8). Operators observe these when they set the basic charge. Price is therefore endogenous. What an instrument must be orthogonal to depends on which components of ξ^C_{kt} the fixed effects remove, so I discuss the two together.

Fixed Effects

Few characteristics of cable systems are observed, so x^C_{kt} consists mainly of fixed effects. The baseline specification includes period effects τ_t and DMA effects. The period effects carry the outside-option index, which rules out MVPD $\times$ period effects: they span the period effects and would absorb λ_t. DMA effects remove persistent local differences in the taste for pay television, over-the-air reception and market structure. What remains in ξ^C_{kt} is a system's deviation from its DMA's average level and from the national path: local changes in channel lineups and promotions, the quality of the operator's broadband service where it is bundled, and operator-specific strategies towards video. Table 6

reports specifications that replace the DMA effects with market-size controls, county-size or census-division trends, or add DMA linear trends.

Standard Instruments

The two standard sources of instruments for price in differentiated-products demand are cost shifters and the prices of the same product in other markets (Hausman 1996; Nevo 2001). Neither suits this setting, for a common reason: operators set prices for regions rather than systems, and negotiate their largest cost, cable network programming, for their footprint as a whole.

Programming fees are the main marginal cost of cable television. Cable networks negotiate them with each operator nationally, so their variation is operator-wide and confounded with the operator's other national choices. Local engineering cost shifters such as plant density do vary across an operator's systems, but cannot move a price set for a region containing many systems. Instruments built from the characteristics of rival products (Berry et al. (1995)) are unavailable for a different reason: few characteristics of cable systems are observed.

Crawford and Yurukoglu (2012) instrument a system's price with the prices of the same operator's systems in other markets. They justify the instrument by the autonomy of local system managers over prices, so that prices across markets share cost shocks but not demand shocks. That premise no longer holds. When an operator sets one rate card for a region, the price of another of its systems in that region is the same decision as the system's own price, and responds to the same regional demand shocks. The variation that survives, across an operator's pricing regions, faces the objection of Bresnahan (1997) with particular force. The dominant demand shock over the sample is the national decline of pay television, which is common to every market an operator serves.

6.3 A Retransmission-Fee Instrument

An instrument that addresses the endogeneity of price, rather than only its measurement, must shift a system's costs and not its quality. One cost is both local and outside the operator's control: the retransmission fees it pays broadcast stations. Each network reaches a DMA through a single affiliate, and a station group negotiates one agreement with each distributor covering all of its stations in the distributor's footprint. Which groups an operator pays for a system's Big Four signals therefore depends on the system's DMA. Two systems of the same operator in the same period face different retransmission costs if the affiliates in their DMAs belong to groups whose fees grew at different rates.

For system k, the instrument is

$$Z_{kt}^{R} = \sum_{m} \omega_{km,0} \sum_{b} n_{mb,0} f_{bt}, \tag{13}$$

where $\omega_{km,0}$ is the system's base-period household share in DMA m, $n_{mb,0}$ is group b's base-period Big Four stations owned in that DMA and f_{bt} is its national retransmission revenue per equivalent affiliate-subscriber-month. The rate paths come from the station groups' financial filings.

The construction separates two sources of variation. The shift f_{bt} is a group's realised bargaining outcome across its footprint: it rises when the group consolidates, when a cycle of its agreements renews, or when it takes a harder line in negotiations. The share $n_{mb,0}$ is a feature of the local broadcast market fixed before the sample begins. Of two otherwise identical systems of the same operator, the one whose local affiliates belong to a group with fast-growing fees faces a steeper cost path, for reasons that have nothing to do with the system itself.

The cost reaches the price through the broadcast surcharge that most operators itemise to recover retransmission fees. The surcharge is set for each pricing region and appears in gross receipts but often not in the posted basic rate, which is one reason the price is receipts per subscriber. Relevance therefore requires operators to pass different cost paths through to different regions. An operator with a single national rate card can contribute to the first stage only through differences in the composition of its subscribers across regions.

Construction

The groups are Nexstar, Gray, TEGNA, Scripps and Sinclair. For systems present at the start of the panel, I freeze DMA household weights in the base year; systems entering later use their first observed footprint. Station ownership and Big Four affiliation exposure are fixed at the start of the panel. The station roster retains one affiliate per DMA and network, avoiding duplicate signals and satellite repeaters.

The rate numerator is each group's annual retransmission revenue from its financial filings. I divide it by twelve times its annual average subscriber base across affiliated stations, expressed in equivalent Big Four units. The denominator follows the group's contemporaneous station portfolio and includes the other network affiliations covered by its revenue, with fixed relative affiliation weights. The exposure in Equation 13 remains restricted to Big Four affiliates. This separates the coverage needed to measure a group's rate from the local stations used to construct a system's exposure.

Exclusion

The economic case for the instrument is that a higher fee for an unchanged broadcast signal raises the distributor's cost without changing what the household receives. This distinguishes it from instruments built from the operator's prices or revenue elsewhere, which move with its demand and with the value of its menu of tiers.

Following Goldsmith-Pinkham et al. (2020), I treat the base-period exposures as the source of identifying variation and the group rate paths as weights. Write $e_{kb} = \sum_m \omega_{km,0} n_{mb,0}$. A sufficient exclusion condition is that each e_{kb} is orthogonal to ξ^C_{kt} conditional on the period effects and on the geographic effects or controls of the specification; in the baseline these are DMA effects. Because DMA effects remove a market's persistent level, the condition concerns changes: which group held a network's affiliation in a DMA at the start of the sample must be unrelated to later changes in the appeal of the cable systems serving that DMA, relative to the national path.

The ownership map is predetermined with respect to the cable side of the market. It is the cumulative outcome of licence transfers and FCC ownership rules over decades, in which cable operators play no part, and no group in the roster holds an interest in a cable system[6]. Predetermination does not by itself establish the condition, however, because the map may be correlated with market characteristics that shaped the subsequent decline of pay television. Three such channels are relevant.

The first runs through operator footprints, which are spatially correlated with those of station groups. Operators concentrated in small markets, such as Cable One, Mediacom and the members of the National Cable Television Cooperative, are the operators most exposed to the groups that consolidated in small and mid-sized markets. The same operators made strategic decisions about video during the sample that are unrelated to any negotiation: Cable One deliberately de-emphasised video from 2013, letting its subscriber base lapse while repositioning around broadband. Such a decision is an operator-wide shock to ξ^C_{kt} that is correlated with exposure through the geography of the footprint. Period and DMA effects leave it in the residual and would attribute the resulting subscriber loss to retransmission costs. MVPD $\times$ period effects would remove it, but they would also absorb the outside-option index.

The second is market size. The consolidating groups are concentrated in small and mid-sized markets, where cord-cutting followed a different path from the largest DMAs. An operator whose footprint spans markets of different sizes is exposed to this within its own systems, so even MVPD $\times$ period effects would not remove it. DMA effects remove differences in the level of demand across markets of different sizes but not differences in

[6]The roster excludes the networks themselves, which own stations in the largest DMAs. Some networks are vertically integrated with distributors, for example NBC with Comcast.

its path. Table 6 therefore reports specifications that control for market size, that replace the DMA effects with Nielsen county-size class or census-division trends, and that add DMA linear trends.

The third is total exposure. For an exposure-weighted mean rate $\bar{f}_t$,

$$Z_{kt}^{R} = \left(\sum_{b} e_{kb} \right) \bar{f}_t + \sum_{b} e_{kb} (f_{bt} - \bar{f}_t). \tag{14}$$

Because fees rose throughout the sample, $\bar{f}_t$ is far from zero, and the first term is total exposure interacted with a trend. Since almost ever DMA has one affiliate of the Big Four networks, the number of Big Four stations is fixed per DMA and so this institutional setup is equivalent to imposing the sum of shares control of Borusyak et al. (2025).

Finally, f_{bt} is a national rate rather than one that leaves out the system's own market. A leave-out rate would have to be built from station-level fee estimates, which are imputed, whereas the groups' filed totals are measured. The national rate is also robust to the demand shocks that most concern the leave-out. It is a ratio of revenue to subscribers, so a fall in subscriptions in a DMA lowers both and leaves the rate largely unchanged, and any single DMA is a small part of a group's footprint.

6.4 Estimation and Inference

For given θ, let $\delta(\theta)$ be the mean utilities satisfying Equation 11 and Equation 12. Let $\xi^{C}(\theta, \alpha^{C}, \beta^{C})$ be the residuals of inner IV-GMM of δ on system characteristics, instrumenting for price. The moment vector $h(\theta, \alpha^{C}, \beta^{C})$ stacks the instrumental-variable moments $\mathbb{E}[Z_{kt}\xi_{kt}^{C}] = 0$ and the differences between predicted and observed survey composition moments. The GMM estimator solves

$$\min_{\theta, \alpha^{C}, \beta^{C}} \quad h(\theta, \alpha^{C}, \beta^{C})' W h(\theta, \alpha^{C}, \beta^{C}), \tag{15}$$

where W is positive definite and the aggregation restrictions enter through $\delta(\theta)$.

I compute this estimator by choosing the mean utilities jointly with the parameters, subject to the aggregation restrictions, following Dubé et al. (2012). Overlapping footprints couple the mean utilities across systems, so nested inversion would require solving a large system at every trial parameter vector. The constrained formulation avoids this inner loop and exploits the fact that each cell contains only a few alternatives. The sparse derivatives and weighting matrices are described in Online Appendix D.

I use two-step GMM. Standard errors come from the GMM sandwich for all parameters jointly: the mean utilities move with the nonlinear parameters through the aggregation restrictions, so the uncertainty in the nonlinear parameters carries through to α and the period effects. The covariance of the instrumental-variable moments is

clustered by system, to allow for persistent unobserved quality within a system. Standard errors for λ_t/α use the delta method on the joint covariance of λ_t and α.

7 Demand Estimates

The baseline estimate of price sensitivity is $\alpha = 0.05526$, with an average own-price elasticity of −1.41. The estimated outside-option path in Figure 7 reaches a gain of $39.7 a month by the second half of 2025. I first report the specification comparisons, then the path of the outside option and household heterogeneity. All dollar amounts are in 2015 dollars per month. Table 6 varies the controls and geographic effects; the row with DMA effects is the baseline.

Across the specifications, the price coefficient lies between 0.0470 and 0.0578. The corresponding average own-price elasticities range from −1.47 to −1.20. The 2025 outside-option gain ranges from $39 to $47 a month. Adding the log share of subscribers buying basic alone, which controls for the menu value of Equation 8, and plant density to the market-size control leaves the price coefficient unchanged.

TABLE 6. Price coefficient and the value of the outside option. Every row includes period effects and the full set of demographic interactions; the row with DMA effects is the baseline. Point estimates are two-step GMM, estimated jointly with the nonlinear parameters; standard errors in parentheses are from the joint sandwich, clustered by system. The outside option is λ_t/α in the second half of 2025, in 2015 dollars a month, with a delta-method standard error. Log households is the log of a system's households in the base period and county-size trends are Nielsen county-size class trends. Elasticities are averages over system-periods.

Absorbed	F	α	Elasticity	Outside option
None	280	0.0525	−1.34	$41.2
		(0.0057)		(5.9)
Log households	203	0.0470	−1.20	$47.3
		(0.0058)		(7.5)
Log households, tier mix and density	215	0.0473	−1.20	$47.0
		(0.0059)		(7.6)
County-size trends	230	0.0470	−1.20	$46.8
		(0.0062)		(8.0)
Census-division trends	222	0.0517	−1.31	$41.5
		(0.0056)		(6.0)
DMA effects	230	0.0553	−1.41	$39.7
		(0.0071)		(5.5)
DMA linear trends	258	0.0578	−1.47	$38.7
		(0.0067)		(4.8)

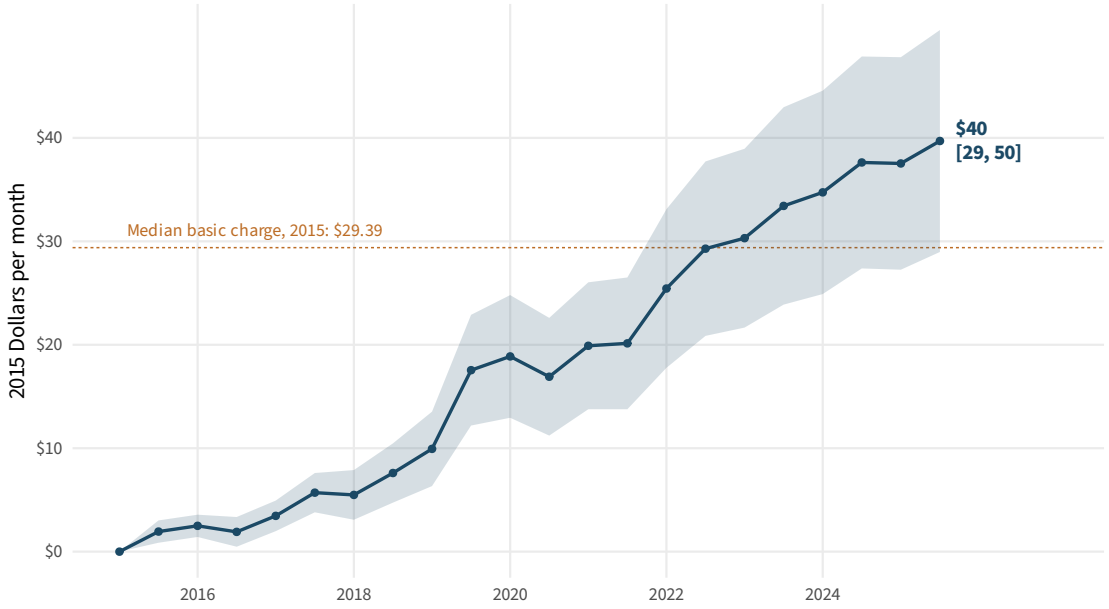

FIGURE 7. Value of the outside option relative to traditional pay television, λ_t/α, in 2015 dollars a month relative to the first half of 2015. The band is a 95 per cent confidence interval, clustered by system, in the baseline specification of Table 6. The dotted line is the median basic charge in 2015, measured by receipts per subscriber.

7.1 The Outside Option

Figure 7 plots λ_t/α for each half-year. The value of the outside option relative to traditional pay television rises from zero in the first half of 2015 to $39.7 a month in the second half of 2025, with a 95 per cent confidence interval of [$29.0, $50.4]. By 2023 the gain exceeds the median basic charge in 2015. The path rises steadily over the decade, with small reversals in 2016, 2018 and 2020. As discussed above, this trend can be interpreted as the monetary value of the relative decline in traditional TV quality vs the outside option.

7.2 Heterogeneity

Table 7 reports the demographic interactions in dollars, for a household one standard deviation above the mean in each characteristic. Age is the margin of price sensitivity. An older household is 1.46 times as price sensitive as the average household and, though imprecisely, values both pay television and the outside option less. Income operates through the outside option instead: a richer household values the outside option $41.4 a month more per standard deviation of local broadband availability, values pay television $8.4 more and is no less price sensitive. Owner-occupation raises the value of pay television by $12.8 a month per standard deviation and lowers price sensitivity, though

TABLE 7. Demographic interactions for a household one standard deviation above the mean in each characteristic, holding the others at their means. The first two columns are in 2015 dollars a month at $\alpha = 0.0553$, the baseline of Table 6. The broadband-access column is η^D times the standard deviation of local broadband availability (0.157); for the average household, η is worth \$47 a month per standard deviation ($t = 5.4$). The price coefficient is α for that household. Absolute t-statistics for the interaction, from the joint covariance clustered by system, are in parentheses.

	Pay television	Broadband access, η^D	Price coefficient
Age	−\$4.1	−\$11.0	0.0806
	(1.4)	(1.2)	(4.2)
Log income	\$8.4	\$41.4	0.0566
	(3.2)	(7.5)	(0.2)
Owner-occupier	\$12.8	\$7.2	0.0418
	(6.1)	(0.9)	(2.3)

imprecisely. Across households, a one-standard-deviation increase in local broadband availability raises the value of the outside option by \$47 a month. Of these gradients, those in age are the most robust. The income and tenure terms change sign under alternative treatments of the broadband term and of the fixed effects that leave the age terms unchanged, so I do not rely on them below.

8 Counterfactuals

I use the estimates to answer three questions. How large would pay television be in 2025 had the outside option not improved after 2015? How much of the decline in subscriptions is due to rising prices rather than the improving outside option? How has the improving outside option changed the elasticity of demand facing distributors and the destination of the subscribers they lose?

8.1 No Improvement after 2015

The first counterfactual replaces the outside-option index in 2025 with its 2015 value, a reduction of 2.09 utils or \$36 a month. Table 8 reports the results. Pay television would reach 69.9 per cent of households rather than 32.2 per cent, and 89.7 million subscribers rather than 41.4 million, 117 per cent more. Penetration fell by 47.9 points over the decade, so resetting the outside option alone undoes 78 per cent of the observed fall. The remainder reflects everything else that changed between 2015 and 2025 and is held at its 2025 value here, including prices, the mean utilities of cable and satellite, choice sets and the composition of households. Cable would recover almost all of its 2015 share, whereas satellite would recover only about half of its loss. Satellite's decline therefore

TABLE 8. Pay-television penetration with the outside option held at its 2015 level. Shares are percentages of households. Other cable is the composite of Section 111 filers that Kagan does not attribute to an operator. The counterfactual holds 2025 prices, choice sets and the mean utilities of every alternative fixed.

	Cable	Satellite	Other cable	Pay television	Subscribers (m)
2015, observed	49.4	27.6	3.1	80.2	93.1
2025, observed	22.8	8.7	0.8	32.2	41.4
2025, 2015 outside option	49.1	19.0	1.8	69.9	89.7

owes more to factors other than the outside option, which in the model enter its mean utilities. Because this counterfactual shifts utilities directly, it does not depend on α, which only converts the shift into dollars.

8.2 Prices or the Outside Option

The second counterfactual decomposes the model-implied change in penetration within the 2025 market into a price effect and an outside-option effect. I evaluate penetration at each combination of 2015 and 2025 prices and 2015 and 2025 values of the outside option and allocate the interaction by a Shapley decomposition. The model implies a fall of 41.3 points from these two sources, against the observed fall of 47.9 points. The outside option accounts for 37.1 points, or 90 per cent, and prices for 4.2 points, or 10 per cent. At the restricted specification's own price coefficient of 0.0390 the outside option accounts for 93 per cent. An alternative design on the entry tier alone, with the posted basic rate as the price, gives 91 per cent.

This share and the 78 per cent of Table 8 answer different questions and are not competing estimates. Their numerators nearly coincide: the outside option accounts for 37.1 points here and 37.6 points there. They differ in the denominator. The first counterfactual divides by the whole observed fall, which also reflects changes in cable and satellite quality, choice sets and household composition. The decomposition divides only by the part of the fall that prices and the outside option jointly imply, and asks how that part splits between them.

8.3 The Outside Option as a Competitive Constraint

The third counterfactual measures how the improvement in the outside option has changed the demand that cable systems face. Table 9 reports the subscriber-weighted own-price elasticity of cable systems and the diversion of their marginal subscribers to the outside option, satellite and other cable systems. In 2015, 27 per cent of the subscribers a cable system lost to a price increase went to the outside option; by 2025

TABLE 9. Own-price elasticity and diversion for cable systems, at $\alpha = 0.0575$. Price is the subscriber-weighted basic charge, measured by receipts per subscriber, in 2015 dollars a month; it differs between the two 2025 rows with fixed prices because the subscriber weights change. The elasticity is with respect to a system's own basic charge and is weighted by subscribers. Diversion is the share of the subscribers a system loses to a marginal price increase that move to each alternative.

			Diversion to		
	Price	Elasticity	Outside option	Satellite	Other cable
2015, observed	\$24.85	−0.84	0.268	0.374	0.331
2025, observed	\$30.35	−1.25	0.759	0.119	0.115
2025, 2015 outside option	\$30.67	−0.94	0.453	0.298	0.231
2025, 2015 outside option and prices	\$22.38	−0.59	0.462	0.285	0.235

76 per cent did, and the elasticity increased in magnitude from −0.84 to −1.25. Holding 2025 prices and choice sets fixed, returning the outside option to its 2015 level lowers the elasticity to −0.94 and diversion to the outside option to 45 per cent. The improvement in the outside option therefore raised the elasticity of cable demand by 33 per cent and diversion outside pay television by 31 percentage points. Satellite was the main destination of cable's marginal subscribers in 2015; by 2025 it is the outside option. Rolling prices back to 2015 as well lowers the elasticity further, to −0.59, while leaving diversion nearly unchanged.

The elasticity with respect to the basic charge is below one in absolute value in 2015. This does not imply that distributors price below the profit-maximising level. The basic charge is the price of access to the operator's whole menu: a subscriber also buys higher tiers, premium channels and, often, broadband, so the margin on a subscriber exceeds the basic charge, and an operator maximising profit over the menu sets the basic charge where demand for subscription is less elastic than a single-product firm would. What matters for the questions of this paper is the change. The improving outside option has made cable demand substantially more elastic and has made the outside option, rather than satellite, the principal alternative to which cable loses subscribers. Both changes alter the incentives that govern how distributors respond to increases in programming costs.

9 Conclusion

A rise in the relative attractiveness of the outside option, not prices, explains most of the decline of traditional pay television. Relative to traditional television, the outside option improved by the equivalent of almost \$40 a month between 2015 and 2025. Had it not

improved, 70 per cent of households would subscribe in 2025 rather than 32 per cent, undoing 78 per cent of the observed decline; of the decline that prices and the outside option jointly imply, a decomposition attributes 90 per cent to the outside option and 10 per cent to prices. The improvement has also changed whom cable competes with. The share of a cable system's marginal subscribers who leave pay television altogether rose from 27 to 76 per cent, and satellite is no longer cable's main alternative.

Upstream ownership changes nevertheless continue to affect what distributors charge. Following a broadcast station acquisition, receipts per subscriber rises by approximately 3.6 per cent in acquired markets relative to other markets served by the same distributor, an estimate that is stable across specifications. Markets where the buyer already owned stations show no comparable increase. The response is consistent with distributors passing higher fees for acquired stations through to subscribers in those stations' markets, but the retransmission fees themselves are not observed, so it is not a measured pass-through rate.

A natural extension models retransmission bargaining alongside these demand estimates, allowing fees and prices to respond to the improving outside option. Another models the bundling of video with broadband, which cable distributors also sell and on which streaming depends. The method of matching predicted quantities to totals over overlapping footprints is not specific to television: broadband and hospital services are also sold by providers whose service areas overlap and whose quantities are reported only in aggregate.

These findings bear on the consolidation of station groups. The case that consolidation harms households rests on higher fees reaching their bills, and the event study shows that acquisitions do reach the basic charge that every subscriber pays. But the constraint on cable pricing increasingly comes from the outside option. A consolidated station group bargains over the revenue of a market that is losing households to an outside option that includes services its retransmission consent does not govern.

Bibliography

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics* 25 (2): 242–62. https://doi.org/10.2307/2555829.

Berry, Steven, Amit Gandhi, and Philip Haile. 2013. "Connected Substitutes and Invertibility of Demand." *Econometrica* 81 (5): 2087–111. https://doi.org/10.3982/ECTA10135.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841. https://doi.org/10.2307/2171802.

Berry, Steven, James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112 (1): 68–105. https://doi.org/10.1086/379939.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. "A Practical Guide to Shift-Share Instruments." *Journal of Economic Perspectives* 39 (1): 181–204. https://doi.org/10.1257/jep.20231370.

Chipty, Tasneem, and Christopher M. Snyder. 1999. "The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry." *The Review of Economics and Statistics* 81 (2): 326–40. https://www.jstor.org/stable/2646869.

Conlon, Christopher, and Jeff Gortmaker. 2026. "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP." *Journal of Econometrics* 257 (September): 105926. https://doi.org/10.1016/j.jeconom.2024.105926.

Crawford, Gregory S., and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review* 102 (2): 643–85. https://doi.org/10.1257/aer.102.2.643.

Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu. 2018. "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica* 86 (3): 891–954. https://doi.org/10.3982/ECTA14031.

Crawford, Gregory S., Oleksandr Shcherbakov, and Matthew Shum. 2019. "Quality Overprovision in Cable Television Markets." *American Economic Review* 109 (3): 956–95. https://doi.org/10.1257/aer.20151182.

Dafny, Leemore, Kate Ho, and Robin S. Lee. 2019. "The Price Effects of Cross-market Mergers: Theory and Evidence from the Hospital Industry." *The RAND Journal of Economics* 50 (2): 286–325. https://doi.org/10.1111/1756-2171.12270.

Dubé, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su. 2012. "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation." *Econometrica* 80 (5): 2231–67. https://doi.org/10.3982/ECTA8585.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, And How." *American Economic Review* 110 (8): 2586–624. https://doi.org/10.1257/aer.20181047.

Goolsbee, Austan, and Amil Petrin. 2004. "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV." *Econometrica* 72 (2): 351–81. https://doi.org/10.1111/j.1468-0262.2004.00494.x.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110 (4): 705–29. https://doi.org/10.1086/340779.

Shcherbakov, Oleksandr. 2016. "Measuring Consumer Switching Costs in the Television Industry." *The RAND Journal of Economics* 47 (2): 366–93. https://www.jstor.org/stable/43895649.

Stahl, Jessica Calfee. 2016. "Effects of Deregulation and Consolidation of the Broadcast Television Industry." *American Economic Review* 106 (8): 2185–218. https://doi.org/10.1257/aer.20110948.

Vistnes, Gregory S., and Yianis Sarafidis. 2013. "Cross-Market Hospital Mergers: A Holistic Approach." *Antitrust Law Journal* 79 (1): 253–93. https://www.jstor.org/stable/43486958.