

Econometric Appendix

ESTIMATION OF THE EFFECT OF FWA ENTRY ON CABLE PRICES

Methodology, empirical results, and econometric detail

James E. Prieger

October 8, 2026

© James E. Prieger 2026

This document is the Econometric Appendix for the report *The Impact of 5G Fixed Wireless Access Entry on Competition in the U.S. Residential Broadband Market*.

Contents

1	Sample Construction.....	1
1.1	Broadband prices and FWA entry data.....	1
1.2	Control variables.....	4
1.3	Summary statistics.....	5
2	Callaway–Sant’Anna Estimation for Staggered Difference-in-Differences	7
2.1	The Callaway–Sant’Anna estimator	7
2.2	Why Callaway–Sant’Anna estimation instead of TWFE estimation	9
2.3	How the estimator handles bias from confounders.....	11
2.4	Enforcing an absorbing treatment: data patching and trimming.....	12
2.5	Selection of control variables and identifying assumptions	13
2.6	Results	16
2.6.1	Baseline and dynamic estimates	16
2.6.2	Robustness and alternative treatment	20
2.7	Translating the price effect into dollar terms.....	22
2.7.1	The percentage effect.....	22
2.7.2	The dollar figure	23
2.7.3	Conservative features of the calculation	23
2.7.4	Sensitivity	23
3	Data Caveats.....	24
4	National Pricing Policies and the Local Counterfactual	26
4.1	The estimator already removes uniform pricing	26
4.2	The installed base is priced separately from the current national “best offer”	26
4.3	The outcome variable measures transaction prices, not uniform offers	28
4.4	The posted menu itself varies by market.....	28
4.5	The objection proves too much	28
4.6	Even if the objection is correct, the estimates are conservative	29
	References	30

Tables and Figures

Table A - 1: Summary statistics for the variables in the regressions	5
Table A - 2: Pre-trends tests by control set (standalone broadband, patch set 1)	15

Table A - 3: Callaway–Sant’Anna estimates and robustness (standalone broadband unless noted)	17
Table A - 4: Group-time cell observations and selection behind the main estimate	21
Table A - 5: Sensitivity of the estimated annual savings to the size of the price effect	24
Figure A - 1: Entry timing and cohort definition	9
Figure A - 2: Illustration of the assumption of parallel trends	14
Figure A - 3: Impact of 5G FWA entry on cable prices: Event-study estimates	18

ESTIMATION OF THE EFFECT OF FWA ENTRY ON CABLE PRICES

Methodology, empirical results, and econometric detail

This appendix documents the difference-in-differences estimates behind the price effects discussed in the main report. Section 1 describes the construction of the estimating sample. Section 2 presents the group-time estimates from the dynamic difference-in-differences model of Callaway and Sant’Anna (2021). Section 3 discusses a few caveats, and Section 4 explains why local counterfactual pricing is consistent with apparently national pricing policies.

1 Sample Construction

1.1 Broadband prices and FWA entry data

The estimating sample combines several sources. Cable broadband prices are drawn from Recon Analytics’ *Consumer Wireless & Internet Pulse* surveys of over 231,000 subscribing households, which record the provider, monthly prices consumers pay, and other characteristics of the plan. The price data are aggregated to the ZIP code level.¹ Information on service availability for cable, fiber, and 5G fixed wireless comes from the Broadband Data Collection (BDC) undertaken by the Federal Communications Commission (FCC), in which every provider reports the individual locations at which it claims residential service is available, together with the maximum advertised download and upload speeds offered there. Throughout, a location counts as served by a given technology only where the provider claims residential (or mixed residential and business) service at advertised speeds of at least 100 Mbps downstream and 20 Mbps upstream. The latter restriction effectively limits fixed wireless access (FWA) to 5G technology (although the actual technology employed is not recorded).

The price and service coverage data are observed on different geographies, and reconciling them is the first step in building the panel. Prices are reported for ZIP Code Tabulation Areas (ZCTAs), a Census geography based on ZIP codes. BDC records are tied to individual locations, which were aggregated to the Census tract for purposes of defining an area that FWA providers choose to enter. While any choice of geographical unit for market entry necessarily involves trade-offs, Census tracts represent a geographic scale well-suited to this analysis. Ideally, the unit of

¹ The Consumer Wireless & Internet Pulse is a continuous survey program that has conducted two million telecom surveys since 2022, described by Recon Analytics as the largest and most current syndicated wireless and home-internet dataset available. Quality control is high; Recon reports rejecting unqualified responses at 4–10 times the industry standard rate. Recon made proprietary broadband service prices from 231,109 households with cable internet service available to CTIA for this project in aggregated form (to preserve anonymity of respondents). The unit of aggregation was a ZIP code × cable provider × plan type (bundled vs. standalone broadband service) × pricing type (promotional vs. regular) cell in a quarter year. For use in the estimations the data were further aggregated to the half-year.

observation would be the geographic unit over which providers make the decision to enter. Small units such as Census blocks would show variation of a different kind: raggedness in a service footprint reflecting propagation and siting constraints rather than distinct decisions to enter. Coarser units, such as ZCTAs to match the Recon data, would mask variation in FWA availability within a single unit and potentially combine many separate decisions whether to enter smaller areas. At the tract level, a change in coverage more plausibly corresponds to a single decision to serve the area, while still preserving meaningful differences in availability across areas.

Therefore, estimation is carried out at the tract level, with each tract matched to its dominant ZCTA (the one to which the tract's 2020 Census population was most heavily allocated).² Thus, entry (along with any covariates) is measured at the tract level while prices are measured at the ZIP code level. Tracts in the same ZIP code therefore share price movements; all reported measures of precision are adjusted accordingly, so that confidence in the results is not overstated.³

The unit of observation (the *panel unit* in econometric parlance) is a cell defined along four dimensions: Census tract, cable internet provider, plan type, and pricing type. Cable internet providers are restricted to the three largest cable companies: Charter, Comcast, and Cox. Plan type distinguishes a standalone broadband subscription from a bundle that also includes some combination of video, voice, and mobile service. Pricing type distinguishes a promotional, introductory rate from the standard, ongoing rate. A single ZCTA or tract can therefore contribute several cells in a given period, namely one for each combination of provider, plan type, and pricing type observed there. Except where noted, the estimates below are for standalone broadband plans, because the composition of a bundle is not observed; specifications that also include bundled plans are reported separately.

Observation periods are half-yearly, running from the second half of 2022 through the second half of 2024, so that the panel has five periods.⁴ For this time period, there are 131,873 price surveys aggregated into 20,974 ZCTA $\times$ provider $\times$ plan type $\times$ pricing type $\times$ period observations, of which 16,453 are for the main estimation sample of standalone broadband plans. The ZCTA-level observations are associated with 176,089 tract $\times$ provider $\times$ plan type $\times$ pricing type $\times$ period observations for the estimations. The main estimation sample (standalone plans in tracts that are

² While it might appear more natural to aggregate up to ZCTA, that area is too large for a meaningful binary measurement of FWA entry. By the start of the sample period, almost all ZCTAs had 5G FWA presence somewhere in them.

³ In particular, standard errors are clustered at the ZCTA level throughout. These cluster-robust standard errors account for arbitrary heteroskedasticity across observations and within-cluster correlation over time, reflecting price dependencies among tracts, providers, plan types, and rate structures within the same ZCTA.

⁴ Although later data are available for both prices and entry, after 2024 5G FWA was so ubiquitous that there was no longer meaningful variation in entry across areas.

not already entered at the beginning of the observation period and can be matched to covariates) contains 12,940 observations.⁵

The outcome variable is the natural logarithm of the average monthly plan price in the cell; a specification using the logarithm of the median price is reported alongside it as a check on the influence of outlying prices.⁶ The treatment indicator records entry by the three large national fixed-wireless providers. A tract is coded as entered in a given period if T-Mobile, AT&T, or Verizon reports at that period's end in the BDC data that it can serve at least one location in the tract that a cable provider reported that it can also serve. Entry is thus defined on the overlap of the FWA and cable footprints. FWA service available to non-cable locations might affect an incumbent cable provider's pricing via *potential* competition. However, whether the entrant has reached customers the incumbent could otherwise serve should have a much larger effect. Because the entry indicator depends only on the conduct of the fixed-wireless providers, it does not vary across cable providers, plan types, or pricing types within a tract. In alternative estimations, entry is allowed to be by any 5G FWA provider, not just the three national-scale providers.

Linking prices to the entry data requires a decision, since the former are averages over a period while the latter pertain to a specific date. Prices represent half-yearly averages (or medians), whereas broadband coverage is recorded as a single point-in-time snapshot at the end of each period, leaving the exact temporal precedence between prices and coverage unknown. The timing convention chosen is to attach end-of-period coverage data to the price periods. Thus, a tract first recorded as entered in period t (i.e., no FWA presence in period $t-1$ but present in period t) may have been entered before the prices observed in that period were set, after them, or partway through the period.⁷ This ambiguity about timing contaminates the period of entry (i.e., period 0 after entry) for purposes of determining the causal effect of entry on prices. Therefore, the

⁵ As will be discussed below, the greatest factor in reducing the total available observations to the estimation sample is dropping always-entered tracts.

⁶ In cases in which the exact median reported price within the (ZCTA $\times$ provider $\times$ plan type $\times$ pricing type) aggregated to the half-year period can be known, it is used (roughly 60 percent of observations). This is possible because for the quarterly data, the min, max, mean, median, and std. deviation of prices are known, allowing mathematical recovery of individual prices whenever there are four or fewer respondents in the quarterly cell. In other cases, the medians for the two (ZCTA $\times$ provider $\times$ plan type $\times$ pricing type $\times$ quarter) observations in the Recon survey data are averaged (weighted by the number of survey respondents in the two quarters).

⁷ Under the end-of-period anchoring for the coverage data, a tract that first shows entry in period t could have gained coverage at any point during the period. Because prices and coverage are not observed on a common clock, no choice of when within the period to anchor the coverage snapshot removes the ambiguity about the relationship between when prices were set and when entry occurred. Anchoring coverage data to the beginning of the period would not eliminate the issue; it would simply shift the ambiguous period back. Under beginning-of-period anchoring, a tract coded as entered as of the start of period t could instead have gained coverage at any point during the prior period, $t-1$, and so it would be $t-1$'s price observation that is contaminated.

estimated treatment effect for the entry period is ignored as contaminated.⁸ By construction of the estimator, no other period's treatment effect (e.g., for the first or second period after entry) will be contaminated by the entry period.

1.2 Control variables

Many control variables were also assembled for the geography and time period of the price and entry data. Control variables for most demographic factors were taken from the Census Bureau's American Community Survey (ACS) 5-year estimates. Population-growth estimates for tracts were obtained from the National Cancer Institute's Surveillance, Epidemiology, and End Results (SEER) program. Industry affiliation of workers resident in the local area, which is related to demand for broadband due to work-from-home, was constructed from the Census Bureau's LEHD Origin-Destination Employment Statistics. Receipt of supply-side Federal broadband subsidies was constructed from the Universal Service Administrative Company's (USAC) High Cost program data. Terrain ruggedness, a deployment cost factor, was calculated from Nunn and Puga's (2012) global elevation-ruggedness grid. A network-cost index was constructed from engineering cost estimates produced by CostQuest's Connect America Cost Model.

For purposes of use in the estimations below, the control variables are grouped into increasingly inclusive sets. Although not all sets end up being required to remove bias from the estimations, all are listed here for the sake of completeness.

Set 0: No control variables. Confounding is addressed only by the differencing inherent in the estimator.

Set 1: Market size and demographics: log number of housing units,⁹ log population density, five-year population growth rate, log median age, race/ethnicity shares (Hispanic, Black non-Hispanic, Asian non-Hispanic, and all-other non-Hispanic, with White non-Hispanic as the excluded reference category), and log average household size. Source: ACS.

Set 2: Adds socioeconomic characteristics: educational attainment (share with a high school diploma or GED, some college, or an associate's degree; share with a bachelor's degree; and share with a master's, doctoral, or professional degree, with less than a high school diploma as the excluded reference category), log median household income, log unemployment rate, and log English-language proficiency share. Source: ACS.

Set 3: Adds infrastructure cost and subsidy measures: log share of large multifamily housing (housing units in structures with five or more housing units out of all housing

⁸ Due to how the Callaway–Sant'Anna estimates are constructed, the entry periods cannot simply be dropped before estimation, as they could with TWFE estimation.

⁹ Variables that were significantly right-skewed enter the regressions in natural logarithms instead of levels, to reduce the influence of outlying observations (as is standard econometric practice).

units),¹⁰ CAF/RDOF subsidy exposure,¹¹ and two deployment-cost indices (terrain ruggedness and block-level costs from a broadband cost proxy model).¹² Sources: as noted in the footnotes.

Set 4: Adds broadband demand-relevant factors: housing vacancy rates, owner-occupancy rates, and baseline (beginning-of-sample) household broadband adoption and log computer-ownership shares. Source: ACS.

Set 5: Adds baseline (beginning of sample) broadband market structure: the share of cable locations for which fiber internet is also available from the same provider, the average number of fiber providers per Census block in the tract, and whether the cable provider in the tract also offers fiber.¹³ Source: author computations from FCC BDC.

1.3 Summary statistics

Summary statistics for the variables used in the regressions are shown in Table A - 1. Where the regressor is in logs, the same variable is also shown in levels for reference. All control sets are included in the table, although not all are used in the preferred estimation (the determination of which is discussed in Section 2.5 below).

TABLE A - 1: SUMMARY STATISTICS FOR THE VARIABLES IN THE REGRESSIONS

Variable	Mean	Std. dev.	Minimum	Maximum
Mean monthly cable price	105.580	62.336	1.670	500.000
Log mean monthly price	4.511	0.537	0.513	6.215
Housing units in tract	1,743.7	712.447	1.000	12,980.0

¹⁰ Multifamily housing share is related to broadband deployment cost because a single run of plant reaches many units in an apartment building versus one in a detached home, lowering the cost per home passed. In addition, multifamily dwellings often permit simpler wiring and bulk building-level access agreements.

¹¹ Federal broadband infrastructure subsidies are included via three variables: the share of housing units with cumulative CAF-subsidized broadband deployment (to any provider) through period t , a binary indicator variable for Cox Communications having received any CAF or RDOF subsidy for any location in the tract and c = Cox, and a similarly constructed variable for Charter. Comcast received no such subsidies during this period. Data on federal subsidies is from USAC.

¹² The terrain ruggedness index is a tract-level measure of how uneven the local terrain is, computed by averaging a fine-grained global elevation-variation grid (Nunn & Puga, 2012) over all grid cells falling within each tract. The other cost index is a tract-level measure of the average per-block cost of deploying broadband service, drawn from CostQuest Associates' Price Cap Block Cost Analysis (built on the FCC's Connect America Cost Model), aggregated from the underlying Census-block cost estimates. While the CostQuest data were publicly available for a time, they no longer are, and the author thanks Dr. Kenneth Flamm for sharing them. Refer to Flamm and Varas (2022) for discussion of these data.

¹³ As elsewhere, all variables mentioned here are restricted to residential broadband service. The distinction between the first and third variables is that the first measures the tract-level share of cable locations upgraded to same-provider fiber (a continuous variable aggregated across all cable providers in the tract), while the third is a binary indicator specific to a single cable provider (the provider to which the price pertains for the observation).

Variable	Mean	Std. dev.	Minimum	Maximum
Log housing units	7.376	0.462	0.000	9.471
Population density	7,297.0	13,609.3	0.053	310,137.5
Log population density	8.061	1.428	-2.937	12.645
Annualized 5-year population growth rate	0.004	0.015	-1.068	0.140
Median age	38.865	7.939	5.700	85.100
Log median age	3.640	0.200	1.740	4.444
% Hispanic	0.230	0.243	0.000	1.000
% Black, non-Hispanic	0.134	0.200	0.000	1.000
% Asian, non-Hispanic	0.076	0.119	0.000	1.000
% all other non-Hispanic (except White)	0.046	0.039	0.000	1.000
Mean household size	2.643	0.564	1.020	18.000
Log mean household size	0.950	0.210	0.020	2.890
% with high school degree	0.529	0.154	0.000	1.000
% with a bachelor's degree	0.215	0.110	0.000	1.000
% with a graduate or professional degree	0.136	0.109	0.000	1.000
Median household income	83,266.9	39,253.9	2,499.0	250,001.0
Log median household income, if available	11.174	0.880	0.000	12.429
Income not available (indicator variable)	0.004	0.066	0.000	1.000
Unemployment rate	0.058	0.045	0.000	0.882
Log unemployment rate	-3.210	1.134	-8.926	-0.125
% speaking English less than very well	0.059	0.085	0.000	1.000
Log % speaking English less than very well	-4.445	2.572	-8.905	0.000
% housing units in 5+ unit structures	0.216	0.241	0.000	1.000
Log % housing units in 5+ unit structures	-3.072	2.835	-9.078	0.000
% housing units with CAF-subsidized broadband	0.005	0.045	0.000	1.000
% housing units with RDOF-subsidized broadband	0.000	0.012	0.000	1.000
Cox received subsidies and obs. is for Cox*	0.000	0.000	0.000	0.000
Charter received subsidies and obs. is for Charter	0.001	0.032	0.000	1.000
Terrain ruggedness index	0.313	0.410	0.000	7.565
Log terrain ruggedness index	-2.009	1.898	-10.069	2.024
CostQuest broadband cost estimate (\$/location)	21.82	9.120	11.71	338.59
Log broadband cost estimate, if available	3.013	0.375	0.000	5.825
Cost estimate not available (indicator variable)	0.005	0.072	0.000	1.000
% housing units vacant	0.087	0.088	0.000	0.857
Log % housing units vacant	-3.233	1.861	-8.564	-0.154
% occupied units owner-occupied	0.621	0.244	0.000	1.000
% households with broadband subscription	0.892	0.085	0.141	1.000
% households with no computer	0.055	0.054	0.000	0.830
Log % households with no computer	-3.794	1.966	-8.702	-0.187
% cable footprint also served by fiber	0.006	0.062	0.000	1.000

Variable	Mean	Std. dev.	Minimum	Maximum
Average fiber providers per block	0.397	0.439	0.000	2.483
Cable provider also offers fiber (indicator variable)	0.048	0.215	0.000	1.000

*Variable is zero everywhere in the estimation sample. Cox Cable received subsidies in 2H2024, but that variable is matched to price data in 1H2025, which is outside of the window for the regressions.

Notes: In estimations using control sets 1–4, $N = 12,940$ observations. With control set 5, $N = 12,918$. CAF is Connect America Fund; RDOF is Rural Digital Opportunity Fund. See text and accompanying footnotes in Section 1.2 for variable definitions.

2 Callaway–Sant’Anna Estimation for Staggered Difference-in-Differences

The preferred estimates come from the group-time approach of Callaway and Sant’Anna (2021), an estimator built for settings like this one, in which different units are treated at different times (what econometricians call *staggered treatment assignment*). The object of estimation is the average effect of 5G FWA entry by the large national providers on the price of cable internet service in the areas they actually entered: the *average treatment effect on the treated* (ATT). The treatment effect estimated here should not be confused with year-on-year changes in overall broadband prices, which are differenced out of the estimation. Rather, it is the difference between (i) cable prices in a tract once 5G FWA has entered and (ii) the cable prices that would have prevailed in that same tract, at the same time, had FWA not entered, holding all else constant. In this way, it isolates the price effect attributable to FWA competition.

2.1 The Callaway–Sant’Anna estimator

The outcome variable for the estimations is the natural log of the average monthly plan price in the ZCTA associated with the tract. Define *entry cohort* g to be the set of tracts in which 5G FWA entry occurred in period g , as shown in Figure A - 1. The Callaway–Sant’Anna estimator computes a separate treatment effect for each entry cohort g and period t , denoted $ATT(g,t)$. Thus, each cell in Figure A - 1 is allowed to have its own treatment effect. For example, $ATT(2,3)$ is the treatment effect in period 3 (2H2023) of entry on prices for Cohort 2, for which FWA entry occurred in period 2 (1H2023); this first period after entry for Cohort 2 is labeled +1 in Figure A - 1. Making the treatment effect specific to a period allows the pricing impact to evolve after entry. For example, the impact on cable prices may grow over time after entry. Each $ATT(g,t)$ compares the observed outcome for units in tracts first treated in period g with the counterfactual outcome that would have prevailed in the absence of treatment, using units in not-yet-treated tracts as the comparison group. This choice of comparison group is important because conventional TWFE can instead use already-treated units as controls for units that enter treatment later. When treatment effects persist, those already-treated units no longer represent untreated outcomes and can therefore contaminate the comparison. The individual $ATT(g,t)$ are then aggregated into two summaries: an overall statistic for the price effect, and a dynamic profile of effects by periods since entry, called an *event study*.

The unit of observation is the tract $\times$ cable provider $\times$ plan type $\times$ pricing type cell discussed in Section 1.1. The Callaway and Sant’Anna (2021) estimator identifies each relevant $ATT(g,t)$ with a two-group, two-period difference-in-differences (DiD) regression. Each 2×2 comparison contrasts the change in outcomes for an entry cohort g against the change for a suitable comparison group (experiencing no entry yet) over the same two periods. For both groups, the change over time is measured relative to the last pre-treatment period (i.e., the last period before entry, the cells labeled -1 in Figure A - 1). Working in changes removes each group’s own price level. Differencing the two groups’ changes then removes any price movement common to both over the same interval. In the absence of covariate adjustment, this is equivalent to a two-way fixed-effects regression. With covariates (a case further discussed in the next subsection), the estimator combines outcome regression with inverse-probability weighting, providing protection against misspecification of either component. Hence, this method is termed *doubly robust*.¹⁴ If covariates are included, they enter measured at the base period (the period just before treatment, $t = g - 1$), so that the comparison remains a clean pre/post contrast not contaminated by covariate changes that could themselves be affected by treatment.

The cohort- and period-specific treatment effects are then aggregated to obtain the overall average treatment effect and dynamic event-study estimates reported below. Identification is based on a conditional parallel-trends assumption, explained in Section 2.5 below. Given the DiD setup, at least one period before and one period after entry are required, and so cohorts 2 through 4 are the only usable ones for the Callaway–Sant’Anna estimations (hence the faded rows in Figure A - 1 for cohorts 1 and 5). The faded rows warrant a clarification, since both cohorts do appear in the tabulation of group-time cells below. Cohorts 2 through 4 are the only cohorts that contribute *post-entry* treatment effects. Cohort 5, entered in the final period, still contributes pre-treatment $ATT(g,t)$, which enter the pre-trends tests and the pre-entry points of the event study. Because Cohort 5 remains untreated until the last period, its units serve as not-yet-treated controls for the earlier cohorts throughout. Cohort 1, entered in the first period, has no pre-treatment period at all and can serve neither as a treatment cohort nor as a control.

¹⁴ For the result that the estimator is consistent if either the outcome or treatment-assignment model is correctly specified, see Sant’Anna and Zhao (2020).

FIGURE A - 1: ENTRY TIMING AND COHORT DEFINITION

Cohort 5	−4	−3	−2	−1	entry
Cohort 4	−3	−2	−1	entry	+1
Cohort 3	−2	−1	entry	+1	+2
Cohort 2	−1	entry	+1	+2	+3
Cohort 1	entry	+1	+2	+3	+4
	2H2022 (Period 1)	1H2023 (Period 2)	2H2023 (Period 3)	1H2024 (Period 4)	2H2024 (Period 5)
	Not yet entered	Entry period (event time 0)	Periods after entry (event time = 1, 2, ...)		

Figure notes: Numbers in cells before and after entry are event times ($t - g$). Post-entry treatment effects are measured against event time -1 , the last period before entry. Pre-entry effects are estimated differently: each compares a period with the one immediately preceding it, so event time -1 is itself an estimated cell wherever a cohort has two or more pre-entry periods, and Cohort 2 contributes no pre-entry cell because its event time -1 is the first period of the panel and has no predecessor to difference against. Event time 0, the period of entry, is included in estimation but ignored in the reported post-treatment ATTs. Cohort 1 is excluded since it has no pre-entry observations. Cohort 5 contributes no post-entry treatment effects, because its only post-entry observation is the period of entry itself; it nonetheless contributes pre-entry cells, and its units serve as not-yet-treated controls for the earlier cohorts. Not shown are observations that are always treated, for which FWA entry occurred before 2H2022.

2.2 Why Callaway–Sant’Anna estimation instead of TWFE estimation

The more familiar two-way fixed-effects (TWFE) regression is not used here, because it is ill-suited to settings in which units are treated at different times. Conventional TWFE estimation can use units that have already received treatment as controls for units treated later, even though the earlier treatment may already have affected their outcomes, a property known as the *forbidden-comparison problem*.¹⁵ This would be akin to a medical study of pharmaceutical efficacy using

¹⁵ In a staggered-entry setting, a TWFE regression implicitly uses already-treated tracts as part of the control group for later-treated tracts. Because those already-treated tracts’ prices have already responded to their own FWA entry, they are poor stand-ins for the untreated counterfactual, and such comparisons are properly regarded as forbidden (Goodman-Bacon, 2021). This is harmless only if the treatment effect is constant

patients who had already received the medicine as a control group for patients treated later, even though the medicine may already have affected the former group's outcomes.

The Callaway–Sant’Anna estimator avoids this problem by constructing each $ATT(g,t)$ using units that have not yet received treatment, including those never receiving treatment, as the comparison group.¹⁶ Thus, for a tract entering treatment in period g , its post-entry outcomes are compared with outcomes for otherwise comparable tracts that remain untreated through the relevant period. Already-treated tracts are not used as controls for later-treated tracts. This preserves the interpretation of the comparison group as representing the untreated counterfactual instead of a mixture of untreated and previously treated outcomes.

The distinction matters empirically in this application because fixed wireless by T-Mobile and Verizon was already broadly deployed early in the sample period. Consequently, a large share of tracts are already treated when the panel opens. A conventional TWFE regression therefore places substantial weight on comparisons in which these already-treated tracts serve as controls for later entrants. Because their cable prices may already have responded to FWA entry, these comparisons can substantially distort the pooled TWFE coefficient by biasing it toward zero.

Empirically, the forbidden-comparison problem (using already-treated units as controls) would be severe here had TWFE estimation been used instead of Callaway–Sant’Anna. Nearly all (97.9 percent) of the weight behind the TWFE estimate comes from comparisons between always-treated tracts and tracts treated later; these comparisons yield an effect of essentially zero.¹⁷ By contrast, comparisons using only never-treated units as controls account for just 0.1 percent of the weight, yet yield an effect of -0.124 , close to the Callaway–Sant’Anna estimate presented below.¹⁸ A complementary diagnostic finds that about 65 percent of the 115,371 underlying group-time

across cohorts and over time; here, the price response evolves with time since entry and plausibly differs across entry cohorts, so these forbidden comparisons can bias TWFE estimates substantially. The Callaway–Sant’Anna estimator avoids the problem by construction, since its component two-period comparisons never use already-treated units as controls.

¹⁶ In studies with many never-treated units, researchers sometimes restrict the control group to those units. In the present context, however, relatively few areas never experience FWA entry, while many more are eventually treated but remain untreated at the time of a given comparison.

¹⁷ The decomposition does not include any control variables and is implemented via a user-written Stata command (Goodman-Bacon et al., 2019). It is computed on the balanced subsample of tracts observed in all five periods, as required by Goodman-Bacon (2021), since the underlying weights are derived under the assumption of a balanced panel. This is less than half of the entire sample, since many ZIP codes do not appear in all survey waves.

¹⁸ The remaining 2 percent of the comparisons are between the groups treated during the sample period. Some of the control groups in this set of comparisons are valid, since they are not yet treated. However, some have already been treated by the focal period, making them problematic for identifying the treatment effect because they have already been exposed to FWA. The estimate for this group as a whole is also near zero.

treatment effects that would enter a TWFE specification carry a negative weight.¹⁹ The negative weights sum to -0.894 , compared with $+1.894$ for the positive weights. Such heavy negative weighting means the overall TWFE estimate need not lie between the smallest and largest of the underlying effects, as a properly convex average would, and in principle could carry the opposite sign of every one of them if treatment effects are heterogeneous (de Chaisemartin and D'Haultfœuille, 2020). These odd properties of TWFE estimation are what motivated the development of the Callaway–Sant’Anna estimator used here. By restricting the comparison group for each $ATT(g,t)$ to units that have not yet been treated, the Callaway–Sant’Anna approach avoids these contaminated comparisons and permits the estimated effect to vary both across treatment cohorts and over time since entry.

2.3 How the estimator handles bias from confounders

Because each cohort and event-time treatment effect is identified from a two-period DiD, the types of confounding factor that Callaway–Sant’Anna estimation removes automatically are straightforward to enumerate. Differencing the outcome within each cell between the base period and the focal period eliminates any time-invariant factors specific to that cell, whether associated with the ZIP code area, provider, plan type, pricing type, or any combination of these. For example, bias from all of the following unobserved factors is removed:

- Any persistent cost or demand factors specific to the local area, including all average characteristics of the Census tract over this time span, whether observed or unobserved.²⁰
- Persistent differences among the cable internet providers in general pricing conventions or strategy (e.g., persistently higher or lower markups or company-wide pricing policies, although not differences in how providers respond to competitive entry).
- Systematic differences during this time in customer demand for the cable company’s service.
- Persistent differences in prices for a single cable provider among its different service areas.
- Average characteristics of the internet plans in the ZIP code over this time span (average speeds, other dimensions of quality such as latency, contract length, etc.) that make them inherently priced above or below plans in other areas, regardless of competition.
- For promotional pricing, the average terms tied to the promotional mix in the local area over this time span (e.g., typical discount depth and length of discounted pricing).
- Local regulatory or franchise conditions specific to a provider and a location (franchise fees, right-of-way costs, terms negotiated with the local franchising authority) that do not change during the sample period (or the time-average of conditions that do change).
- Average network vintage or technology generation for a cable-provider’s plant in the area (e.g., predominant DOCSIS version, average age of physical infrastructure) that shapes

¹⁹ This diagnostic is based on the method of de Chaisemartin and D’Haultfœuille (2020) applied to the full sample (there is no requirement for a balanced sample).

²⁰ By definition, the average-over-time of any variable is time-invariant (fixed in the sample period), and hence removed by differencing over time.

both the provider's underlying cost structure there during the sample period and, potentially, customer demand (via the perceived quality of service).

- Any other idiosyncratic factor, not captured by the categories above, that does not change between the base period and the focal period for that particular combination of area, cable provider, plan type, and pricing type.

Any time-invariant factor specific to a given cell is therefore differenced out, once the explicitly included control variables (those discussed in Section 1.2) have been accounted for. This holds whether the factor is fixed at the cell level or simply averaged over the sample span. As a result, such factors cannot bias the estimated impact of FWA entry.

The DiD comparison also differences out shocks common to all cells in a given period. Thus, any unobserved factor that affects prices across all markets in a particular period is removed by the time differencing. Examples include:

- National movement in the cost of inputs such as labor, capital, and materials.
- Secular trends in the price of residential broadband service.
- Nationwide macroeconomic conditions (e.g., inflation, interest rates, or aggregate demand shifts) affecting pricing power or consumer willingness to pay across all markets simultaneously.
- Industry-wide strategic or regulatory developments in a given period that affect all cable providers' pricing at the same time. Examples include policy changes at the FCC, any activity in national spectrum auctions during the relevant half-year, coordinated industry responses to broader competitive pressure, or changes in the average state policy environment during the period.

Just as differencing over time removes cell-specific factors, differencing treatment units against control units in the same period ensures that any remaining unobserved factors common across all markets in a given period will not bias the estimated effect of FWA entry.

2.4 Enforcing an absorbing treatment: data patching and trimming

The Callaway–Sant'Anna estimator, like most staggered-adoption estimators, requires treatment to be an *absorbing* state: once there is 5G FWA entry in a tract, that tract must count as treated in every subsequent period. In the raw dataset assembled from the BDC data, this does not always hold: in a minority of cases (fewer than 10 percent), a tract shows no reported FWA coverage in a period despite coverage having previously been reported in the same tract. Most often this apparent “exit” is only for one period (over 77 percent of such cases). Given the nature of investment in broadband networks, it is likely that such entry gaps are spurious for the large national FWA providers and merely reflect reporting omissions.

The no-exit condition (i.e., the absorbing-state requirement) is enforced by patching and trimming the series. In the first approach (termed *patch 1* below), once entry is first observed in a tract, the tract is coded as treated in all later periods. Patch 1 repairs only single periods of apparent exit, however, which are most plausibly data errors. Once the patch would have altered two or more

consecutive observations in a tract, all later observations for that tract are dropped.²¹ Under this restriction, fewer than 1 percent of observations are dropped. As a check on this choice, estimates are also reported for the case of *patch 0*, which makes no such repairs and instead drops the offending observations; as shown below, the two approaches yield essentially the same estimate.


2.5 Selection of control variables and identifying assumptions

The estimator identifies each $ATT(g,t)$ under a *conditional parallel-trends* assumption: absent entry, average prices in a tract with entry would have evolved in parallel with the comparison group, after conditioning on the control variables X . Figure A - 2 illustrates the parallel trends assumption (where the plotted trends are interpreted as having the effects of the control variables already netted out). This assumption is not directly testable, but it implies that pre-entry trends should already be parallel, which is testable through testing of the pre-treatment $ATT(g,t)$. In the pre-entry periods (those cells labeled with negative numbers in Figure A - 1), since no treatment has yet occurred, the pre-treatment $ATT(g,t)$ cannot reflect any effect of FWA entry; it can only capture a difference in trends between the eventual treatment group and its comparison group that predates entry. Under parallel trends that difference should be zero, so a nonzero estimate would signal that the two groups were already drifting apart for reasons unrelated to entry, which would undermine the case that they would have continued on parallel paths after entry.²² Because parallel trends are assumed to hold *conditional* on covariates, the choice of controls directly determines the estimate of what would have happened to cable prices without FWA entry. Following recent econometric guidance (Callaway and Sant'Anna, 2021; Roth et al., 2023), blocks of control variables capturing key market characteristics are included, and joint pre-trend tests are reported to check whether any pre-treatment divergence is detectable.

²¹ In other words, only a single period of missing presence is recoded to “entered”; runs of two or more consecutive missing periods are instead trimmed from the dataset. For example, if the raw pattern for FWA presence across periods in a tract is 0-0-1-1-0-1, patch 1 recodes it to 0-0-1-1-1-1. Apparent exits at the end of the sample period are treated the same way: 0-1-1-1-1-0 becomes 0-1-1-1-1-1. Where trimming applies, observations before the first trimmed period are retained: a raw pattern of 0-1-1-0-0-1 becomes 0-1-1-*m-m-m*, with *m* denoting a missing value.

²² Passing the test that the pre-treatment effects are individually and jointly zero is necessary but not sufficient for conditional parallel trends. Indeed, no test can be sufficient, because conditional parallel trends is a claim about the counterfactual path the treated tracts would have followed after entry, and that path is never observed. Instead, the test looks for observable implications that would occur if the assumption were false, so passing it means only that no such evidence was found.

FIGURE A - 2: ILLUSTRATION OF THE ASSUMPTION OF PARALLEL TRENDS



Notes: The figure provides a stylized illustration of the conditional parallel trends assumption. In particular, it depicts the implication that, absent treatment, the difference in the outcome between the treatment cohort and the comparison group would have remained constant over time (conditional on the covariates included in the regression). The dashed red line represents this counterfactual outcome path for the treatment cohort in the absence of entry. The treatment effect at each post-entry period, $ATT(g, t)$, is the vertical distance between the observed outcome for the treatment cohort (solid red line) and its counterfactual outcome (dashed red line). The counterfactual outcome is identified from the change in the comparison group's outcome between the last pre-treatment period ($g - 1$) and period t .

Table A - 2 reports the pre-trends test results across control sets for the standalone-broadband sample (bundles excluded). With no controls the joint test is borderline, and adding the demographic block alone (control variables Set 1 as defined in Section 1.2) makes it worse. Adding the socioeconomic block (Set 2) brings the test well above conventional significance thresholds ($p = 0.185$ for the joint test and none of the individual pre-treatment periods shows anything close to rejecting the hypothesis that the treatment effect is zero). Control set 2 is therefore adopted as the working specification for the standalone broadband estimation; the block with added cost variables (Set 3), for which the pre-trends pass more comfortably still ($p = 0.438$), is reported as a robustness check.²³ The remaining control sets are not employed in this estimation.

²³ For the broader sample that also includes bundled plans, control set 2 does not pass the pre-trends test: $p = 0.078$, while set 3 does ($p = 0.281$ from the joint test, and the smallest p -value from the individual pre-treatment periods is 0.454; neither comes close to rejecting that there are no pre-treatment effects). Therefore, in the estimation including bundles reported in Table A - 3, the latter control set is used.

TABLE A - 2: PRE-TRENDS TESTS BY CONTROL SET (STANDALONE BROADBAND, PATCH SET 1)

Control set	Joint χ^2 test statistic	p-value of joint test	Largest pre-trend ATT(t)	Smallest p-value of individual tests	Assessment and Use
Set 0	10.56	0.103	-0.027	0.326	Marginal (joint test p is near 10%). Not used
Set 1	11.95	0.063	-0.054	0.289	Fails (joint test p < 10%). Not used
Set 2	8.80	0.185	-0.040	0.462	Passes. Baseline specification
Set 3	5.87	0.438	-0.030	0.559	Passes. Robustness check

Notes: The dependent variable is $\ln(\text{average monthly price})$. Joint χ^2 test: Tests the null hypothesis that all pre-treatment ATT(g,t) estimates are jointly equal to zero. Largest pre-trend ATT(t): The aggregate pre-treatment event-study coefficient ATT(t) (representing the average treatment effect at relative time t averaged across cohorts) with the largest absolute magnitude. The reported p-value corresponds to the minimum p-value among the z-tests of the ATT(t) estimates across the three pre-treatment relative periods; since the std. errors of the ATT(t) estimates differ, the p-value reported may not correspond to the ATT(t) shown. Unlike the estimates reported elsewhere in this section, which use the wild bootstrap, the pre-trend tests are computed from asymptotic variance estimates. The bootstrap procedure returns standard errors and confidence intervals but not p-values, and the joint test across pre-treatment periods requires an estimated covariance matrix. The device used for Table A - 3, in which p-values are bounded by locating zero relative to confidence intervals of varying coverage, serves for individual estimates but does not extend to a joint test.

A related concern governs which control variables belong in this regression. An estimate that survives the inclusion of a long list of control variables looks more convincing than one that does not, so applied work is often drawn toward adding controls until the list is exhausted. However, not all possible covariates belong in a treatment effects regression. Several of the variables in Sets 4 and 5, particularly broadband adoption, computer ownership, fiber overlap, and provider density, are not fixed background features of a market but are themselves plausibly outcomes of FWA entry. Entry can change adoption, change concentration, and reshape the very competitive configuration these variables are meant to describe. Conditioning on their *contemporaneous* values would make them “bad controls”,²⁴ partialing out part of the causal channel under study and biasing the estimated treatment effect. This is one of the reasons these sets of controls are not employed here for the preferred estimation.²⁵

The worry is not confined to Sets 4 and 5. Even the demographic and socioeconomic variables in Sets 1 and 2 may not be entirely pre-determined. Suppose higher-income, better-educated households sort toward areas with cheaper, faster, or more numerous broadband options, and that those options depend in turn on FWA entry. Income and education, both part of control set 2, would then be picking up part of a sorting process that FWA entry itself sets in motion. The control

²⁴ In the sense of Angrist and Pischke (2009).

²⁵ Although, as discussed above, the main reason they were not used above is that control set 2 satisfied the conditional parallel trends assumption required for identification under the Callaway–Sant’Anna estimator and no further controls were needed.

set would not be composed solely of pre-determined characteristics but would lie downstream of entry, shaped through internet service prices. Including such control variables in a regression means conditioning on variables that lie on the causal path from treatment to outcome. In that case, the estimated treatment effect would capture only the portion of entry's total effect on price that did not travel through the sorting channel.²⁶ The bias lands on the estimated treatment effect even though the reverse causality runs between the controls and price rather than between entry and price directly.

The Callaway–Sant’Anna estimator addresses this problem directly by conditioning on covariates measured in the *earlier* period of each of the individual two-period DiD comparisons that make up the estimation procedure.²⁷ Evaluating controls prior to entry within each group-time comparison guarantees that time-varying variables do not absorb the downstream impacts of FWA entry.

In addition to the parallel trends assumption, the estimator also requires that entry not be *anticipated*: prices in a tract must not respond to fixed wireless entry before that entry is recorded in the coverage data. Anticipatory pricing is conceivable in this market, since carriers market service in advance of general availability and cable operators may know of competitors’ plans regarding buildout. Two observations bear on the no-anticipation assumption. As just shown, the pre-treatment “treatment effects” are small and statistically indistinguishable from zero, individually and jointly, which is what would be expected were anticipation absent. Second, anticipation works against finding an effect of entry. A cable provider that lowers its price preemptively ahead of entry contaminates that tract-cell for use as a not-yet-treated control for an earlier cohort. Both channels narrow the measured gap between entered and not-yet-entered tracts. The estimates reported here are therefore underestimates with respect to anticipation: if it were present, the true effect of entry would be larger in magnitude than what is reported.

2.6 Results

2.6.1 Baseline and dynamic estimates

The main estimate of the effect of 5G FWA entry by the three large providers (T-Mobile, AT&T, and Verizon) on cable internet prices is **-0.129** (s.e. 0.059, $p < 0.05$; 95% CI [-0.247, -0.010]), corresponding to a reduction of about 12.1 percent in average cable prices over the two half-years

²⁶ In this case, since the treatment effect coefficient on D is left to explain only whatever portion of the price response did not pass through the income/education sorting channel, one may suspect that there is attenuation bias (i.e., bias toward zero in the estimated β). However, depending on the direction of the relationships involved, the bias can be in any direction. See Rosenbaum (1984) for the seminal result and Caetano et al. (2022) for discussion of the bad controls problem in the staggered difference-in-differences setting.

²⁷ Unlike the fixing of certain control variables in control sets 4 and 5, this base period is not the start of the overall sample, but the earlier (pre-treatment) period in each specific two-by-two DiD estimation, as described in Section 2.1.

following entry.²⁸ This is the shaded row (row 5) of Table A - 3. The estimate averages the post-entry ATT(g,t) across the first and second periods after entry and excludes any group-time cell resting on 40 or fewer control observations. The remainder of this subsection explains how that specification was reached, beginning with the unrestricted aggregate over all post-entry periods and showing why it is not the appropriate figure to report.


TABLE A - 3: CALLAWAY–SANT’ANNA ESTIMATES AND ROBUSTNESS (STANDALONE BROADBAND UNLESS NOTED)

Row	Specification	ATT (β)	Asymptotic s.e.	Significance level	Bootstrap 95% CI
1	Overall ATT, all post-entry periods	-0.272	0.063	< 0.001	[-0.394, -0.150]
2	– log median price (std. pricing)	-0.280	0.073	< 0.001	[-0.418, -0.142]
3	– including bundles (control set 3)	-0.192	0.055	< 0.001	[-0.301, -0.083]
4	1 st and 2 nd periods after entry only	-0.133	0.054	$0.01 < p < 0.05$	[-0.241, -0.024]
5	Main specification: event-times 1–2, control group $N_{gt} > 40$	-0.129	0.059	$0.01 < p < 0.05$	[-0.247, -0.010]
6	Robustness: control set 3	-0.133	0.059	$0.01 < p < 0.05$	[-0.252, -0.014]
7	Robustness: control set 4	-0.104	0.049	$0.01 < p < 0.05$	[-0.198, -0.010]
8	Robustness: control set 5	-0.101	0.050	$0.01 < p < 0.05$	[-0.200, -0.002]
9	Robustness: standard pricing only	-0.129	0.060	$0.01 < p < 0.05$	[-0.247, -0.012]
10	Robustness: no patching (patch set 0)	-0.130	0.060	$0.01 < p < 0.05$	[-0.249, -0.012]
11	Alternative treatment: any 5G entry	0.051	0.117	$p > 0.10$	[-0.172, 0.274]

Notes: All specifications are for log average prices, use control set 2, and patch set 1 unless the row header states otherwise. Treatment is FWA entry by T-Mobile, Verizon, or AT&T; the dependent variable is $\ln(\text{mean monthly price})$ except in row 2, where it is $\ln(\text{median monthly price})$. Inference: 3,999 wild-bootstrap replications, clustered at the ZCTA level. Reported p -value intervals are assessed from bootstrap confidence intervals of varying coverage probability; only the 95 percent CI is reported in the table. “Control group $N_{gt} > 40$ ” trims group-time cells with 40 or fewer control observations; applies to rows 5–11. Row 11 has treatment redefined to be 5G FWA entry by any provider, not just T-Mobile, Verizon, and AT&T. The shaded row (5) is the main estimate.

²⁸ In log-linear regressions, exact percentage changes are calculated as $[\exp(\beta) - 1] \times 100$. Here, $\exp(-0.129) - 1 \approx -0.121$, or roughly a 12.1 percent reduction.

FIGURE A - 3: IMPACT OF 5G FWA ENTRY ON CABLE PRICES: EVENT-STUDY ESTIMATES



Notes: Event-study estimates of the effect of 5G FWA entry by T-Mobile, Verizon, and AT&T on log average cable prices (Callaway–Sant’Anna estimation with control set 2, standalone broadband, patch set 1, and no trimming of small control groups). Points are average treatment effects by periods since entry ($ATT(t)$); vertical bars are 95 percent wild-bootstrap confidence intervals (3,999 replications). The estimate at event time -1 is not normalized to zero; event time -1 is an estimated cell (see notes to Figure A - 1).

The unrestricted estimate averages the post-entry treatment effects across all treated tract-periods since entry.²⁹ Under control set 2, this is -0.272 ($p < 0.001$),³⁰ representing a large, highly statistically significant reduction in prices of roughly 24 percent.³¹ However, that figure is driven almost entirely by the most distant post-entry horizon, as the event-study profile in Figure A - 3 makes clear.

The figure shows that pre-treatment estimates are near zero (roughly -0.04 to $+0.02$) and are jointly and individually statistically insignificant, consistent with the pre-trends test results from Table A -

²⁹ The average is computed with each cohort weighted by its share of treated units, as is standard with this estimation method. Estimation is implemented using Stata’s `csdid` command (Rios-Avila et al., 2021). Inference throughout uses 3,999 wild-bootstrap replications, clustered at the ZCTA level.

³⁰ The bootstrap procedure returns a std. error and a 95 percent confidence interval, but not a p -value. Standard statistical significance (5 percent level) can be assessed by seeing whether zero falls outside the 95 percent confidence interval. For this estimate in particular, the 99.9 percent confidence interval does not include zero, hence $p < 0.001$.

³¹ See footnote 28.

2. The entry effect in the period of entry itself (event time 0, where event time is $t - g$) is small (-0.03) and marginally insignificant (zero lies just inside the 95 percent bootstrap confidence interval). Recall from Section 1.1 that this period is to be ignored and is excluded from all estimates of the overall treatment effect.

More importantly, the post-entry effect grows steadily in magnitude: about -0.08 one period after entry, -0.19 two periods after, and a very large -0.56 three periods after entry. The treatment effects for event times 2 and 3 are individually significant at the 5 percent level, while that for event time 1 is significant at only the 10 percent level. Averaging the three effects produces the highly significant overall ATT of -0.272 mentioned above (see row 1 in Table A - 3).

A parallel estimate on the log *median* price gives a similar estimated ATT of -0.280 , significant at the 0.1 percent level and slightly larger than the mean-based estimate (see row 2 in the table), consistent with the main result not being an artifact of a few outlying prices. Expanding the set of prices from the surveys to include bundles (some combination of internet, video, voice, and mobile services) decreases the impact of FWA entry to -0.192 . This muted effect suggests that cable providers chose not to defend their customer base primarily through multi-product bundle discounts, concentrating their largest price reductions on standalone internet instead.³² Both estimates from these variants are for all post-entry-periods, and so are to be compared with the -0.272 reported in row 1.

As anticipated above, Figure A - 3 shows that the unrestricted ATT is driven largely by the third post-entry period, which rests on a single cohort with only 40 control observations. To be conservative, the aggregation is therefore restricted to the first two post-entry periods, where control-group coverage is substantially stronger. This restricted estimate yields an ATT of -0.133 ($p < 0.05$; see row 4), about a 12.5 percent reduction in cable internet prices. To be more conservative still, the main Callaway–Sant’Anna estimate excludes in addition *any* component of the first three post-entry periods that relies on a control group with 40 or fewer observations. This drops not only the third period after entry, as in row 4, but two other comparisons that affect the estimates for the first two periods. This leaves the estimate only slightly changed, at the -0.129 reported at the beginning of this subsection (row 5). Apart from the third period after entry, then, the results do not depend on these comparisons.

Table A - 4 documents exactly which group-time comparisons enter that main estimate. Of the sixteen $ATT(g,t)$ cells the estimator can form, six are pre-period cells, four are the entry period itself, which is excluded, one lies beyond the retained event-time window, and two are removed for not containing more than 40 control observations. This leaves three cells, all at event-times 1 or 2, that are combined into the reported estimate. Because they are aggregated with cohort-share weights that are themselves estimated, the asymptotic standard error shown in Table A - 3 is computed by

³² While the survey data include prices for bundled plans, the exact components of the bundle are unknown. Thus, while the estimation controls for the difference in average prices of bundles vs. standalone broadband internet over time within each tract $\times$ cable provider $\times$ plan type $\times$ pricing type cell, it does not control for other characteristics of the bundle. Hence these data are not explored further.

the delta method using the full covariance between the weights and the ATT coefficients. The confidence intervals are instead computed by the wild bootstrap.³³ The p -value ranges reported in Table A - 3 are based on bootstrap confidence intervals.³⁴

2.6.2 Robustness and alternative treatment

The lower rows of Table A - 3 vary the main specification along several dimensions. Using the richer control set 3 (which also passes the pre-trends tests reported in Table A - 2; see row 6 of Table A - 3) yields an ATT of -0.133 , a bit larger than the main estimate, while control sets 4 and 5 lead to slightly smaller, yet still statistically significant, estimates. The more consequential contrast in Table A - 3, however, lies not among the overall ATT estimates in rows 5 through 8, which vary only from -0.129 to -0.101 , but in whether event time 3 is included in the average effect. Compared to the preferred specification (row 5), including event time 3 (the estimation in row 1)³⁵ more than doubles the estimated effect relative to restricting to event times 1 and 2, as noted above; the same comparison under control set 5 (from a regression not reported in the table) yields a similarly large multiple, about 1.85. Whichever control set is used, then, the event-time window (or, equivalently in this dataset, the dropping of control groups without more than 40 observations), not the choice of controls, is what governs the magnitude of the estimated effect and hence of the cost savings discussed below.

The final robustness checks involve the type of pricing and the entry-gap patching. Restricting to standard (non-promotional) pricing gives an ATT of -0.129 , and turning off the entry-gap patching of the dataset gives -0.130 . Both figures are very close to the main estimate and remain significant at the 5 percent level, so the main result is robust to these alternative specification choices.

³³ The confidence intervals throughout are computed via the multiplier (wild) bootstrap recommended by Callaway and Sant'Anna (2021), with 3,999 replications and weights drawn at the ZCTA level. The multiplier bootstrap perturbs each cluster's contribution to the estimator's influence function by an independent random weight, which preserves the panel structure and the within-ZCTA dependence, accommodates the estimation of the propensity score and outcome regression underlying the doubly robust estimator, and supports simultaneous confidence bands across the event-study coefficients. The weights follow the two-point distribution of Mammen (1993), which provides an asymptotic refinement.

³⁴ If the 95% bootstrap confidence interval excludes zero but the 99% CI admits it, a row is marked as $0.01 < p < 0.05$. For rows marked as $p < 0.001$, 9,999 bootstrap replications were used to more accurately determine the 99.9% CI.

³⁵ Comparing rows 1 and 5 involves two restrictions: is event time 3 included, and are control groups with 40 or fewer observations included. The latter restriction implies the former, but not vice versa, since some of the ATT(g,t) included in the average effects for event times 1 and 2 also lack more than 40 observations, as will be shown in Table A - 4.

TABLE A - 4: GROUP-TIME CELL OBSERVATIONS AND SELECTION BEHIND THE MAIN ESTIMATE

Cohort (entry period)	Period after entry	Observations in the treatment group	Observations in the control group	Status
2	0	4,406	866	Entry period; ignored
2	1	3,994	306	Included in all estimates
2	2	3,458	70	Included in all estimates
2	3	4,150	40	Beyond event-time window; included only in rows 1–3. Would also be excluded by trimming (control $N \leq 40$).
3	–1	498	368	Pre-period
3	0	492	354	Entry period; ignored
3	1	398	80	Included in all estimates
3	2	476	38	Trimmed (control $N \leq 40$); included only in rows 1–4
4	–2	210	158	Pre-period
4	–1	208	146	Pre-period
4	0	192	76	Entry period; ignored
4	1	196	38	Trimmed (control $N \leq 40$); included only in rows 1–4
5	–3	30	128	Pre-period
5	–2	30	116	Pre-period
5	–1	20	56	Pre-period
5	0	40	52	Entry period; ignored

Notes: As explained in the text, the period of entry is not included in the estimates of the ATT due to unclear timing regarding the prices and entry. Pre-period observations are used to produce the pre-entry effects shown in Figure A - 3. The Cohort 5 rows illustrate the point made in Section 2.1 above: Cohort 5 contributes pre-period cells and an entry-period cell that is ignored, but no post-entry treatment effects.

The specifications above restrict the treatment to entry by AT&T, T-Mobile, and Verizon. It is worth asking whether the price response is specific to these national providers, or whether fixed-wireless entry by any company moves incumbent cable prices. There are many smaller and regional wireless internet service providers (WISPs). Indeed, in June 2024, 232 distinct providers offered 5G FWA residential internet service (minimum advertised download speed of 100 Mbps).³⁶ Nevertheless, the lion’s share of 5G FWA coverage is provided by the three largest providers. During the sample period, an average of 33 percent of locations served by any 5G fixed wireless,

³⁶ In that period in the FCC BDC data there are 105 WISPs offering licensed 5G FWA residential internet service (minimum advertised download speed of 100 Mbps), marketed under 118 brands, and 154 WISPs offering 5G residential service under Licensed-by-Rule (LBR) shared-spectrum access, marketed under 163 brands. There are 27 providers that appear in both lists, for a total of 232 distinct providers.

cable, or fiber internet provider was served by one of the three large FWA providers, whereas including all 5G-speed WISPs into the mix increases that coverage to only 39 percent.

To investigate, in row 11 of Table A - 3 treatment is redefined as entry by any 5G FWA provider. Pooling all FWA entrants, the estimated effect is statistically indistinguishable from zero, and the point estimate is no longer negative. The treatment effect of entry by the three major FWA operators is therefore not merely a relabeling of a broader effect from any WISP; when the fringe of smaller providers is folded in, the competitive response by the cable incumbents disappears. The pattern is what market structure would predict. T-Mobile, AT&T, and Verizon market fixed-wireless service as national, mass-market products, with the brand recognition, retail presence, and pricing scale to pose a credible substitute for cable broadband. Most smaller WISPs, whatever their local importance, apparently do not exert the same pressure on an incumbent cable operator's price.

One qualification belongs with this comparison. Redefining the treatment does more than widen the set of entrants. Because entry by any provider occurs earlier, and in more tracts, than major-carrier entry alone, the broader definition reassigns tracts to different entry cohorts and rebuilds each cohort's control group accordingly. Row 11 therefore rests on a different set of group-time comparisons, and part of the contrast with the major-carrier result reflects a composition effect rather than a difference in competitive effect alone. Two considerations nonetheless support the substantive reading. First, an analogous two-way fixed-effects check points the same way while holding the sample nearly fixed.³⁷ Second, the coverage that the broader definition adds is modest; including all 5G-speed WISPs raises coverage from 33 to 39 percent of locations.

2.7 Translating the price effect into dollar terms

The estimated treatment effect is expressed in log points, and the main report translates it into a dollar figure: roughly \$776 million saved on cable internet bills in the first year following entry, for the markets the three large FWA providers entered during 2023. This subsection sets out how the estimate maps into a proportional saving per subscriber, why applying it as a single multiplier is conservative, and how sensitive the dollar figure is to the precision of the estimate.

The DiD coefficient, $\beta = -0.1287$, implies that prices in entered tracts are $\exp(\beta) = 87.9$ percent of their counterfactual level, a reduction of 12.1 percent; equivalently, the counterfactual price is the observed price multiplied by $\exp(-\beta) = 1.137$, a gross-up of 13.7 percent. All these figures recur below.

2.7.1 The percentage effect

The main estimate is the ATT of -0.129 from row 5 of Table A - 3, which averages the treatment effect over the first and second half-years after entry. The log-linear specification implies that this

³⁷ A parallel two-way fixed-effects specification, run separately for major-carrier entry and for entry by any 5G FWA provider, yields the same directional contrast reported here: the price effect is smaller in magnitude when the broader set of entrants is used, while the two samples differ in size by only about 0.7 percent.

corresponds to prices in entered tracts dropping to 87.9 percent of the level they would otherwise have reached, a reduction of 12.1 percent against that counterfactual level.³⁸

The saving per subscriber is the gap between the price paid after entry and the price that would have been paid without it. Measured against the counterfactual price the gap is 12.1 percent of that price; the multiplier on the observed price actually paid after entry to find the counterfactual price is thus 1.137.³⁹ The base price used in the savings calculation is an average of prices observed after entry, so the latter figure is the multiplier used throughout.

2.7.2 The dollar figure

The remaining steps of the calculation (the base monthly price to which the multiplier is applied, the conversion of cable-serviceable locations into cable-subscribing households, and the aggregation over the first twelve months after entry for the two 2023 entry cohorts) are documented in the Consumer Savings Appendix to the main report, where they form Part I of its savings calculation (see Sections 3 and 4 and Table 8 of that appendix). The \$776 million is the identified-cells component of a larger estimate developed in that appendix, which extends the price effect to cable subscribers elsewhere in the FWA footprint and to FWA subscribers themselves, and puts total consumer savings over 2023 through 2025 at \$6.4 billion to \$28.5 billion in current dollars, depending on how far the effect is extrapolated beyond the cells that identify it.

2.7.3 Conservative features of the calculation

The savings calculation applies a single multiplier to every cell, which treats the proportional effect as common across them. That assumption is conservative. Were the effect β to vary across cells (indexed by unit i and period t), the correct aggregate would apply $\exp(-\beta_{it})$ cell by cell. The exponential is convex, so by Jensen's inequality the average of $\exp(-\beta_{it})$ across cells exceeds $\exp(-\bar{\beta})$, the same function evaluated at the average effect. Using a single multiplier built from the aggregate ATT therefore understates the sum, by an amount that grows with the dispersion in effects across cells. The event-study profile in Figure A - 3 indicates that such dispersion is present over event time; accounting for dispersion across i would strengthen this point.

2.7.4 Sensitivity

The dollar figure inherits the statistical imprecision of the ATT. When the price impact of entry is assumed to be 50 percent higher or lower than the main estimate (values that fit comfortably into its 95 percent confidence interval), aggregate annual savings vary from \$1.2 billion down to \$376 million (Table A - 5). The point estimate is distinguishable from zero at the 5 percent level, and the savings figure is a strictly monotone transformation of β ; the savings estimate is therefore

³⁸ See footnote 28.

³⁹ Let p_0 be the counterfactual (higher) price, p_1 be the actual (lower) price. Per footnote 28, $p_1 = \exp(\beta) p_0$. Then expressing the unobserved p_0 in terms of the observed p_1 yields $p_0 = \exp(-\beta) p_1 = m p_1$. Here, with $\beta = -0.1287$, the multiplier m is 1.137.

significant as well.⁴⁰ Thus, while the qualitative finding that entry lowered cable prices stands, the dollar magnitude is bounded only loosely.

TABLE A - 5: SENSITIVITY OF THE ESTIMATED ANNUAL SAVINGS TO THE SIZE OF THE PRICE EFFECT

Case	β	Multiplier, $\exp(-\beta)$	Total annual savings (\$ millions)
Half the main estimate	-0.0644	1.0665	375.7
Main estimate	-0.1287	1.1373	776.4
Half again as large	-0.1931	1.2129	1,203.7

Notes: β is the Callaway–Sant’Anna estimate from row 5 of Table A - 3. The outer rows scale it by 0.5 and by 1.5. Both lie inside the 95 percent wild-bootstrap confidence interval for β , reaching a little over halfway to each bound. The multiplier converts the price actually paid after entry into the counterfactual price. The base prices and household counts are held at the values used in Part I of the Consumer Savings Appendix (Table 8 of that appendix), and the shaded row is the reported case, which reproduces the Part I total there.

3 Data Caveats

The greatest limitation of the study is that the first usable entry period in the dataset is the first half of 2023, after much FWA entry had already occurred. Indeed, 5G FWA service was already available in 81 percent of Census tracts by the end of 2022 (although service was available to only 32 percent of broadband-serviceable locations).⁴¹ However, no comparable price data are available from an earlier period, making it impossible to extend the analysis further back in time.

The estimations above do not control for possible differences in advertised plan download and upload speeds between control and treatment groups.⁴² The Recon Analytics household survey underlying the price data also asks households what maximum speed their broadband plan offers. However, a large share of respondents (37 percent in the present sample) answer that they do not know, so maximum speed is missing for a meaningful share of underlying observations. A separate issue is that in the aggregated data provided by Recon for this project, only the modal (i.e., most common) reported speed is available for each cell. This creates a mismatch between the price data, which reflect a mean or median across all households in a cell, and the speed data, which reflect only the single most common response. That mismatch precludes constructing an accurate

⁴⁰ This statement assumes that estimation error in the ATT is the only error involved in the savings estimate.

⁴¹ The former statistic is the author’s calculation from the analysis of the FCC BDC data for residential service for the mainland U.S. The latter statistic is from the Internet Access Services report (status as of December 31, 2022; FCC 2024a), Figure 30, which includes all serviceable locations and all states. The main difference between the two figures, however, is that the former measures the share of Census tracts with any 5G FWA availability at all, whereas the latter measures the share of individual broadband-serviceable locations that could obtain the service.

⁴² However, differencing over time removes permanent differences in speeds between the two groups, and differencing between groups removes any trend in speeds common to both.

speed-adjusted price measure at the cell level, such as average (or median) price per unit of speed, a measure sometimes used in the empirical broadband literature. Constructing an ad hoc measure such as average price divided by modal speed is ill-defined, because the mode, mean, and median summarize a cell's distribution in different ways, and the plans driving the modal speed statistic may not be the same plans driving the average or median price.

Notwithstanding this limitation, modal plan speed can be tested directly as an outcome. If so, then the reported price effects may in part be capturing differences in broadband quality. To investigate, the trimmed Callaway–Sant’Anna exercise was repeated de novo with the natural log of modal speed (where known) as the outcome, using the 7,393 observations for which a modal speed is known (compared to the 12,940 observations available for the price estimates).⁴³ The result is a negative but statistically insignificant estimated treatment effect, with a 95 percent confidence interval of $[-0.585, 0.268]$ that comes nowhere close to excluding zero. The estimate is consistent with two readings: that entry produced no speed response, in which case no quality adjustment to the price effects is called for; or that a true speed effect exists but is small enough that it cannot be distinguished from random sampling variation. Under the second reading, the price effects reported above would overstate the quality-adjusted price reduction.

Taken together, these potential limitations suggest that the price results, while robust to the several specification checks reported above, would benefit from corroboration once data become available that record speed at the same granularity as price and that resolve “don’t know” responses through follow-up or imputation.

A final point is more a study design choice than a limitation (and may be of interest only to econometricians). The Recon survey is conducted for business analytics and therefore follows common industry practices for recruiting participants into its panel.⁴⁴ Consequently, it is not a probability sample: households are not selected through a known random mechanism with computable inclusion probabilities, so design-based standard errors cannot be constructed.⁴⁵ The clustered standard errors reported throughout this appendix are therefore *model-based*, treating the observed data as realizations of an underlying stochastic process.⁴⁶ This is the appropriate framework here, since the object of interest is the causal relationship between FWA entry and prices rather than a description of the households surveyed. It does mean, though, that the reported standard errors and confidence intervals quantify sampling uncertainty under the assumed model, not uncertainty arising from how the survey recruits participants, and they do not

⁴³ The control set was selected via the same pretrend-testing procedure described in Section 2.5 above (which here selects control set 0).

⁴⁴ Such practices typically involve recruiting respondents by advertising on the internet and offering small incentives for participation.

⁴⁵ This is a routine feature of commercial survey data, given the practical constraints and intended business purposes of such data collection.

⁴⁶ Under the finite-population approach, by contrast, the target of estimation is a summary of U.S. broadband households that currently exist, and the variability of the estimates arises only from repeated probability sampling from that finite population.

correct for any lack of representativeness beyond that addressed through the included covariates (e.g., the demographic controls in control set 2).

4 National Pricing Policies and the Local Counterfactual

A reader may object that the cable operators studied here do not set prices locally. Charter, Comcast, and Cox each typically advertise a standard set of rates at any one time, as reflected in the broadband consumer labels (colloquially known as the “nutrition labels”) filed with the FCC, and advertise promotional offers that are, on their face, the same across the country. If the price of a given plan does not vary across a provider’s footprint, the objection runs, then no local price response to local entry is possible, and the counterfactual price that the estimator constructs for an individual Census tract is not a meaningful object. Because the objection goes to the validity of the entire approach, it is addressed here at length.

4.1 The estimator already removes uniform pricing

Begin with a matter of construction. The component of price that the objection describes is precisely the component that the Callaway–Sant’Anna estimator differences away before anything is estimated. A pricing policy that is uniform across a provider’s footprint and unchanging over the sample period is absorbed entirely by differencing over time, as noted among the confounding factors listed in Section 2.3. A pricing policy that is uniform across the footprint but moves over time (a national rate increase, a nationwide refresh of promotional offers, an industry-wide response to margin pressure) is absorbed by differencing between units. What remains, and what alone can identify the treatment effect, is the geographically differentiated residual: the part of the price paid that moves in some tracts and not in others.

The objection therefore leads to a prediction that the treatment effect should be zero. That prediction is testable, and the estimates reject it at conventional significance levels across every robustness specification reported in which treatment is major-carrier entry. Whatever the published rates may indicate, the prices that households report actually paying moved differentially with FWA entry, with no comparable movement beforehand.

This argument establishes that prices move differentially with entry; it does not by itself establish that the moving component is a price cut at a given speed. If entry induces some customers to switch to cheaper, slower cable plans rather than inducing providers to lower any given plan’s price, the cell-average price could fall through a compositional shift in the plan mix rather than through a price response. As noted in Section 3, however, a direct test finds no statistically significant effect of entry on modal plan speed. This result does not support the alternative interpretation that the price decline reflects a shift toward lower-speed plans.

4.2 The installed base is priced separately from the current national “best offer”

There is a mechanism through which such differential movement can arise without any change to the published price. National uniform prices and an advertised promotion for the current set of

best offers govern customer *acquisition*. The great majority of a cable operator's subscribers, however, are its installed base, which has historically been priced through an entirely different process. A promotional rate expires; the subscriber rolls to the standard rate; and whatever reduction the subscriber obtains after that requires contacting the provider. The size of the concession granted at that point is discretionary and generally unadvertised. It is determined case by case by a representative who can observe what alternatives are available at the customer's service address. The concession may take the form of a service credit, a waived equipment or installation charge, an extension of promotional pricing, a speed upgrade at no additional cost, or discounts offered via bundling other services. Nothing in a national set of posted prices constrains how liberally that discretion is exercised in a market a major FWA provider has just entered relative to one it has not. Differential retention pricing of this kind generates exactly the tract-level variation in prices paid that the design exploits, while leaving the posted price untouched.

A second mechanism also leads to differential pricing apart from a national uniform standard. "Tier flattening" is a well-known phenomenon in the broadband industry in which consumers who remain on legacy plans with slower speeds often pay nearly as much as consumers on newer plans offering higher speeds, reflecting consumer inertia and the failure of some customers to switch to better-value offerings (Consumer Reports, 2022). Thus, the prevalence of legacy-plan customers in a local area will affect average prices. This customer mix is partly endogenous to the cable company's pricing strategy, because the firm can choose whether and how aggressively to migrate legacy customers to newer plans or retain them on older, higher-priced tiers.

The installed base is an active margin, and the operators' own conduct shows that they manage it separately from the offers aimed at acquiring new customers. In September 2024, the final period of the estimation panel, Charter introduced its *Life Unlimited* pricing across its 41-state service area. The company automatically raised the speeds of existing Spectrum Internet subscribers at no additional charge and moved to all-in bundled prices carrying guarantees of up to three years (Charter Communications, 2024). In April 2025, Comcast announced a five-year price guarantee on Xfinity internet, the longest it had offered on residential broadband. That guarantee went to new customers first, with existing customers permitted to *repackage* into it thereafter (Baumgartner, 2025) — an extension to the installed base that required the subscriber to elect it. Charter, seven months earlier and facing the same competitive pressure, had done the opposite, moving its existing subscribers itself.

Thus, the two firms differ on the length of the commitment and on how it reaches existing subscribers. Charter has kept a three-year guarantee rather than matching Comcast's five-year term.⁴⁷ What the two operators disagree on is the profit-maximizing duration of the commitment

⁴⁷ Asked about the difference in April 2026, Charter's chief executive said the company had run trials of five-year guarantees and price locks without seeing the lift it was looking for and that Charter would adopt the practice if it saw the approach working elsewhere (Neenan, 2026). He has separately described multi-year guarantees as a sound competitive response that he expects to keep.

and how it reaches existing subscribers. Neither pattern of conduct fits a picture of mechanically uniform pricing.

4.3 The outcome variable measures transaction prices, not uniform offers

It matters what the outcome variable actually measures. The prices used in these estimations are the amounts that surveyed households report paying, not posted offers. The transaction price is the list price net of a substantial and locally administered wedge: retention credits, waived fees, promotional extensions, loyalty adjustments, bundle discounts, and equipment charges that are levied in some cases and forgiven in others. Consumer Reports' review of more than 22,000 consumer broadband bills documents how far the amount billed may have departed from any posted price. For some providers, more than half the bills in the sample carried a discount; for others, none did. Discounts typically ran \$10 to \$50. More than a dozen providers levied company-imposed fees at their own discretion (Consumer Reports, 2022). Whatever uniformity exists on the national offers need not carry through to the bills households actually receive, and it is the bill that this analysis takes as its outcome.

4.4 The posted menu itself varies by market

Even the posted menu of prices is less uniform than the objection supposes. Comcast introduced its *NOW* brand in a limited set of markets in April 2024 with prepaid, no-contract service at \$30 per month for 100 Mbps and \$45 for 200 Mbps. Only over the following weeks were the offers extended across the provider's service area, and availability continues to be determined at the service address (Comcast, 2024). Charter's headline broadband prices under *Life Unlimited*, \$30 per month for 500 Mbps and \$40 for 1 Gbps, are conditional on the customer's also taking two lines of Spectrum Mobile or Spectrum video service (Charter Communications, 2024). Thus, the broadband price a household actually pays depends on its mobile purchase, which varies with local mobile market conditions.⁴⁸ Charter has also marketed a guarantee of \$1,000 in first-year savings to households that move their mobile service away from AT&T, T-Mobile, or Verizon. Prices on any given broadband "nutrition label" filed with the FCC can hold constant while the menu offered at a given address, and the price ultimately realized on it, differ across tracts.

4.5 The objection proves too much

The objection, if accepted, would require setting aside a substantial body of prior evidence on cable companies' targeted responses to entry. In the older literature on cable video, Goolsbee and Petrin (2004) estimated that prices would have been roughly 15 percent higher without direct broadcast satellite entry; this is a counterfactual price of exactly the kind at issue here. Chu (2010) found that satellite entry generally induced cable operators to raise quality and lower price, with the direction of the response varying by market. The Commission's own price surveys report the same pattern descriptively: as of January 1, 2010, expanded basic cable video prices averaged 1.5

⁴⁸ Note, however, that bundled pricing such as this is excluded from the main estimation sample (but see row 3 of Table A - 3).

percent lower for incumbent cable operators in communities with a rival operator, and 9.6 percent lower for the rivals themselves (Federal Communications Commission, 2012).

In the broadband era, the competitive-response literature has concentrated on the quality margin, given the paucity of publicly available data on prices. Wallsten and Mallahan (2010) and Molnár and Savage (2017) both find higher delivered speeds in markets with more providers. Prieger, Savage, and Molnár (2014) estimate quality reaction functions for incumbent aDSL providers across roughly a thousand local California markets and find higher delivered speeds where a gigabit fiber competitor is present; this is a competitive response identified at the local market level, within a single state. For the national-pricing objection to hold, 5G fixed wireless would have to be the one entrant to which cable's documented tendency for local competitive response does not extend. This would run counter to the literature documenting such responses to wireline overbuilders, to satellite, and to fiber, and to the fact that the cable systems it studied have since been consolidated in large part into the three operators examined here.

4.6 Even if the objection is correct, the estimates are conservative

The final point is that no version of the objection can inflate the estimated effect. The objection admits two forms, and they fail in different directions: the strong form is a hypothesis the data reject, and the weak form would result in an attenuation that leaves the reported estimates conservative. Consider first the strong form, in which the prices households actually pay for service with particular speeds are uniform across a provider's footprint. As noted in Section 4.1 above, the estimates in Section 2 reject that supposition.

A weaker form of the objection presumes that the cable provider has two responses to entry: a proximate response that lowers prices in the entered markets, and a general response that lowers prices across its entire footprint. The latter price change will be a function of the footprint-weighted average of local market conditions. Entry by a fixed wireless provider in some tracts shifts that average toward more competitive conditions. Through the general response, prices therefore fall throughout the footprint, including in tracts where no entry occurred. Tracts assigned to the comparison group are then partially treated: they receive the weaker general response but not the stronger proximate one. The estimator recovers the proximate response alone, and the contrast between entered and not-yet-entered tracts is compressed by exactly the general response. The estimated treatment effect therefore understates the total reduction in prices attributable to entry.

The reported estimates are therefore a lower bound on the benefit to consumers. This parallels the argument made in Section 2.5 regarding anticipation: to the extent that national pricing spreads a local response across untreated areas, it attenuates the measured gap.

References

- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press. <https://doi.org/10.1515/9781400829828>
- Baumgartner, J. (2025, April 15). Comcast unleashes five-year price lock on broadband. *Light Reading*. <https://www.lightreading.com/cable-technology/comcast-unleashes-five-year-price-lock-on-broadband>
- Caetano, C., Callaway, B., Payne, S., & Rodrigues, H. S. (2022). *Difference in differences with time-varying covariates* (arXiv:2202.02903). arXiv. <https://arxiv.org/abs/2202.02903>
- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
- Charter Communications. (2024, September 16). *Spectrum announces unprecedented customer commitment, free Internet speed lifts and new bundled pricing* [Press release]. PR Newswire. <https://www.prnewswire.com/news-releases/spectrum-announces-unprecedented-customer-commitment-free-internet-speed-lifts-and-new-bundled-pricing-302249060.html>
- Chu, C. S. (2010). The effect of satellite entry on cable television prices and product quality. *The RAND Journal of Economics*, 41(4), 730–764. <https://doi.org/10.1111/j.1756-2171.2010.00119.x>
- Comcast. (2024, April 16). *Comcast launches NOW—a new brand and product portfolio that redefines low-cost internet, mobile and TV* [Press release]. <https://corporate.comcast.com/press/releases/comcast-launches-now-low-cost-internet-mobile-and-tv>
- Consumer Reports. (2022, November 17). *Broadband pricing: What Consumer Reports learned from 22,000 internet bills*. <https://innovation.consumerreports.org/FINAL.report-broadband.november-17-2022-2.pdf>
- de Chaisemartin, C., & D'Haultfoeulle, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996. <https://doi.org/10.1257/aer.20181169>
- Federal Communications Commission. (2012). *Report on cable industry prices* (DA 12-377). <https://docs.fcc.gov/public/attachments/DA-12-377A1.pdf>
- Federal Communications Commission. (2024a). *Internet access services: Status as of December 31, 2022*. <https://docs.fcc.gov/public/attachments/DOC-405486A1.pdf>
- Federal Communications Commission. (2024b). *Internet access services: Status as of December 31, 2023*. <https://docs.fcc.gov/public/attachments/DOC-405488A1.pdf>
- Federal Communications Commission. (2025). *Internet access services: Status as of June 30, 2024*. <https://docs.fcc.gov/public/attachments/DOC-411463A1.pdf>
- Flamm, K., & Varas, P. (2022). Effects of market structure on broadband quality in local U.S. residential service markets. *Journal of Information Policy*, 12, 234–280. <https://doi.org/10.5325/jinfopoli.12.2022.0006>
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277. <https://doi.org/10.1016/j.jeconom.2021.03.014>
- Goodman-Bacon, A., Goldring, T., & Nichols, A. (2019). *bacondecomp: Stata module for decomposing difference-in-differences estimation with variation in treatment timing* (Version 1.0.5) [Computer

- software]. Statistical Software Components S458676, Boston College Department of Economics. <https://ideas.repec.org/c/boc/bocode/s458676.html>
- Goolsbee, A., & Petrin, A. (2004). The consumer gains from direct broadcast satellites and the competition with cable TV. *Econometrica*, 72(2), 351–381. <https://doi.org/10.1111/j.1468-0262.2004.00494.x>
- Mammen, E. (1993). Bootstrap and wild bootstrap for high dimensional linear models. *The Annals of Statistics*, 21(1), 255–285. <https://doi.org/10.1214/aos/1176349025>
- Molnár, G., & Savage, S. J. (2017). Market structure and broadband Internet quality. *The Journal of Industrial Economics*, 65(1), 73–104. <https://doi.org/10.1111/joie.12106>
- Neenan, J. (2026, April 24). Charter down 120,000 broadband subs. *Broadband Breakfast*. <https://broadbandbreakfast.com/charter-down-120-000-broadband-subs/>
- Nunn, N., & Puga, D. (2012). Ruggedness: The blessing of bad geography in Africa. *Review of Economics and Statistics*, 94(1), 20–36. https://doi.org/10.1162/REST_a_00161
- Prieger, J. E., Savage, S. J., & Molnár, G. (2014). *Quality competition in the broadband service provision industry* [Paper presentation]. 42nd Research Conference on Communication, Information and Internet Policy (TPRC), Arlington, VA, United States. <https://ssrn.com/abstract=2418676>
- Rios-Avila, F., Sant’Anna, P. H. C., & Callaway, B. (2021). *CSDID: Stata module for the estimation of difference-in-difference models with multiple time periods* (Statistical Software Components S458976) [Computer software]. Boston College Department of Economics. <https://ideas.repec.org/c/boc/bocode/s458976.html>
- Rosenbaum, P. R. (1984). The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society: Series A (General)*, 147(5), 656–666. <https://doi.org/10.2307/2981697>
- Roth, J., Sant’Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2), 2218–2244. <https://doi.org/10.1016/j.jeconom.2023.03.008>
- Sant’Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1), 101–122. <https://doi.org/10.1016/j.jeconom.2020.06.003>
- U.S. Census Bureau. (2024). *Types of computers and internet subscriptions* (American Community Survey 1-Year Estimates, Subject Table S2801). <https://data.census.gov/table/ACSST1Y2023.S2801?q=broadband+s2801&g=010XX00US&y=2023>
- Wallsten, S., & Mallahan, C. (2010). *Residential broadband competition in the United States* (SSRN Working Paper No. 1684236). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1684236