

Online Appendix for "Uniform vs Competing Standards: A
Structural Analysis of the U.S. Wireless Telecommunications
Industry"*

Michela Bonani[†]

September 4, 2026

*

[†]Department of Economics and Business, Universitat Pompeu Fabra, P.O. Ramon Trias Fargas 25-27, 08005
Barcelona, Spain; michela.bonani@upf.edu

A Appendix: Data Construction and Variable Definitions

This appendix provides a detailed description of the data collection, spatial mapping, and variable construction used to develop the estimation database. The data pipeline combines multiple raw sources at differing levels of geographic granularity, including individual households, census blocks, and Designated Market Areas (DMAs).

To aggregate block-level indicators and individual households to the DMA level, I use mapping tables provided by the U.S. Census Bureau. The Census Bureau assigns numeric GEOIDs that uniquely identify all statistical geographic areas. Relying on these GEOIDs and federal FIPS codes, I crosswalk the appropriate demographic data from censuses and surveys to the relevant market levels for data analysis, interpretation, and mapping.

A.1 Market Selection and Choice Sets

To construct a reliable panel of competitive interactions suitable for structural estimation, I apply a filtering process to the raw data. A DMA-year is retained in the final estimation sample based on the following criteria:

1. **Data Availability and Dynamic Choice Sets:** Coverage data for the four nationwide carriers must be observable in the FCC database. However, firms are not strictly required to be active in every market. If a provider has zero native coverage and zero roaming coverage in a given DMA, it is economically inactive in that location. In such cases, its market share is set to zero, and the provider is excluded from the consumer's choice set for that specific market.
2. **Lagged State Availability:** The preceding year's coverage data ($t-1$) must be available to serve as the baseline state for current investment decisions. Because the model relies on changes in coverage, the 2014 coverage data is strictly used to initialize the 2015 investment state.
3. **Non-Negative Investment:** Neither compatible firm within a pair can exhibit a substantial, unexplained contraction in total coverage. Small negative fluctuations (below a tight percentage-point threshold) are treated as measurement noise in the FCC data and zeroed out. However, massive network contractions are excluded, as the structural model strictly evaluates the decision to expand or maintain network infrastructure (i.e., investment $y \geq 0$). Note that markets where firms choose exactly zero investment are intentionally kept, as they identify baseline costs and represent valid corner-solution equilibrium outcomes.

Applying these rules yields a final unbalanced panel of 145 unique DMAs observed across the 2015–2018 period.

A.2 Consumer Panel and Market Shares

Consumer choices are extracted from the NielsenIQ Consumer Panel. I map observed brand and product descriptions to the four parent carriers (AT&T, Verizon, T-Mobile, and Sprint). Mobile Virtual Network Operators (MVNOs) and regional carriers are classified as the outside option.

Because the NielsenIQ data relies on survey sampling—which can under-represent certain demographic cells in smaller DMAs—I use observed national market shares to adjust and validate the local DMA-level market shares. Households are assigned to DMAs based on their reported county FIPS codes using mapping files provided by NielsenIQ and the Census Bureau.

A.3 Demographics and Market Size at the DMA Level

Household-level demographics are directly retrieved from the NielsenIQ Panel as explained in the main text. To construct population-level distributions for the estimation routine, I use the American Community Survey (ACS) 5-year estimates.

I first define the age and income distribution for each DMA-year by leveraging the ACS data. Notably, the survey reports age in 9 bracketed categories. I take the midpoint of each bracket and fit a normal distribution to estimate the market-specific mean (μ) and standard deviation (σ) of age for each DMA-year. Mirroring the age definition, household income brackets are used to construct income quartiles.

Lastly, total market size is defined as the total ACS-estimated population residing in the DMA.

A.4 Network Coverage and Geographic Aggregation

The primary source for infrastructure data is the FCC Form 477. Carriers report coverage differently over time, requiring two distinct processing approaches.

For the **2014 Baseline Coverage** (which serves as the lagged state X_{t-1} for 2015 investment decisions), I extract the 2014 FCC coverage shapefiles. I use Census block boundary files to extract the centroid coordinates (latitude and longitude) of every census block. Using a point-in-polygon ray-casting algorithm, I determine whether a block’s centroid falls within a provider’s technology-specific coverage polygon (CDMA EV-DO for Verizon/Sprint; WCDMA/HSPA for AT&T/T-Mobile).

For the **2015–2018 Panel**, the FCC provides block-level CSVs detailing the deployment of 3G technologies as the percentage of geographic area covered within each block. To define the total coverage percentage for a firm in a given DMA-year pair, I first compute the sum of the land area of all covered blocks divided by the total land area of the DMA. In a second step, I disentangle the covered area to compute the distinct spatial areas of interest.

Specifically, to properly estimate the structural network investment game, total DMA coverage must be decomposed into its strategically relevant sub-components (overlap, unique coverage, catch-up/roaming, and empty space). Because the raw FCC data reports the percentage of area covered within each census block, I must first calculate the spatial intersections at the block level before aggregating to the DMA level.

The exact intra-block topological distribution of base stations is unobserved; however, for the vast majority of census blocks, the spatial intersection is mathematically deterministic. Specifically:

- If at least one firm covers exactly 100% of the block, or if at least one firm covers exactly 0% of the block, the overlap, unique, and empty areas can be calculated with absolute certainty.
- It is only in the fractional “in-between” cases, where both compatible firms report partial coverage strictly between 0% and 100%, that the exact topological overlap is unobserved.

For these specific fractional blocks, I apply a maximum spatial correlation assumption. Given the extremely small geographic size of a census block, I assume that if both firms report partial coverage, they are covering the exact same populated areas of that block. Moreover, given the geographical constraints exogenous to the providers and the high degrees of coverage in each block, it is likely that firms build base stations close to each other. Because the percentage of blocks falling into this uncertain, fractional category is remarkably small¹, this assumption aggregates up to highly realistic and reliable spatial metrics at the DMA level.

Once the components are resolved for every census block, they are aggregated to the market level strictly within technology-compatible pairs of firms. Notably, the total 3G coverage, total roaming access, overlapping coverage and empty areas for firm j operating under standard r in

¹AT&T reports a share of 97.8% of blocks being extreme cases, T-Mobile a share of 96.9%, Verizon 97.8%, and Sprint 96.9%, with an overall percentage across all firms of 97.3%.

market m at time t are constructed as follows:

$$\begin{aligned}
x_{jrm}^{3G} &= \frac{\sum_{c \in C_m} Area_{cjrmt}^{3G}}{\sum_{c \in C_m} AreaTotal_{cmt}} \times 100, \\
xRoam_{jrm}^{3G} &= \sum_{\substack{k \in \mathcal{J}_r \\ k \neq j}} \frac{\sum_{c \in C_m} \max(0, Area_{ckrmt}^{3G} - Area_{cjrmt}^{3G})}{\sum_{c \in C_m} AreaTotal_{cmt}} \times 100, \\
xOverlap_{jrm}^{3G} &= \sum_{\substack{k \in \mathcal{J}_r \\ k \neq j}} \frac{\sum_{c \in C_m} \min(Area_{cjrmt}^{3G}, Area_{ckrmt}^{3G})}{\sum_{c \in C_m} AreaTotal_{cmt}} \times 100, \\
xEmpty_{jrm}^{3G} &= \sum_{\substack{k \in \mathcal{J}_r \\ k \neq j}} \frac{\sum_{c \in C_m} (AreaTotal_{cmt} - \max(Area_{cjrmt}^{3G}, Area_{ckrmt}^{3G}))}{\sum_{c \in C_m} AreaTotal_{cmt}} \times 100.
\end{aligned} \tag{1}$$

where $Area_{cjrmt}^{3G}$ denotes the 3G-covered area of firm j in block c within the set of blocks $C_m \in m$, $AreaTotal_{cmt}$ is the total area of block c , $1_{\{Area_{cjrmt}^{3G}=0\}}$ equals 1 if provider j does not have any coverage in the block, and $\mathcal{J}_r$ is the set of firms deploying standard r .

This bottom-up construction perfectly enforces the geographic identities required by the MPEC solver at the market level. For any firm j and its compatible rival k , the data pipeline guarantees the following exact spatial accounting:

$$\begin{aligned}
overlap_pct + my_complementarity_pct_j &= competitor_coverage_pct_k \\
overlap_pct + competitor_complementarity_pct_k &= own_coverage_pct_j \\
total_joint_coverage_pct + pair_empty_pct &= 100\%
\end{aligned}$$

A.5 Spatial Concentration (HHI) and Urban Ratio of Empty Areas

Because land area alone does not dictate the economic value of coverage, I construct a population-weighted Spatial Herfindahl-Hirschman Index (Spatial HHI) to proxy for demand concentration within a DMA's empty area. Using the ACS block-group population crosswalked to census blocks, I compute the variance in population density across the specific geographic blocks comprising the empty area within a DMA-year pair. Moreover, to account for geographic characteristics of the empty area, I compute the Urban Ratio as the fraction of urban blocks over the total number of blocks within a market's empty space.

Specifically, I define a Spatial Herfindahl-Hirschman Index (HHI) of population across the c distinct census blocks within market m 's empty area, defined as $SpatialHHI_{m,empty} = \sum_{c \in C} (pop_c / pop_{m,empty})^2$. Furthermore, I define the population-weighted urban share of the empty area as $UrbanRatio_{m,empty} =$

$\sum_{c \in C} \left(\frac{pop_c}{pop_{m,empty}} \right) \mathbb{I}_{c,urban}$, where $\mathbb{I}_{x,urban}$ indicates whether block c is urban.

These metrics are computed separately for the UMTS and CDMA compatible groups and serve as critical demographic shifters for the spatial matching friction in the network investment model. It should be noted that both measures vary across DMAs and dynamically over years within a DMA, driven by the fact that the empty area boundaries shift over time as firms continually allocate investments to expand their networks.

A.6 Summary Statistics

Table 1 presents the evolution of the 145 markets over the 2015–2018 sample period. A market is included in the sample if it exhibits non-negative investments by active firms and has observable data for both the current and preceding periods. Over this four-year span, the industry exhibits a steady decline in the average consumer price, dropping from \$38.30 in 2015 to \$33.25 in 2018.

Table 1
Summary Statistics - Key Variables (2015-2018)

	2015	2016	2017	2018
Number of Markets (total)	29	49	85	59
Number of Urban Markets (total)	19	4	16	15
Number of Carriers per Market	4.00	4.00	4.00	4.00
Number of Markets per Carrier	29.00	49.00	85.00	59.00
Price (USD)	38.30	35.46	34.38	33.25

Note: This table shows key variables for the 145 markets.

Table 2 details the distributions of spatial coverage, revealing an asymmetry between the industry leaders and the followers. Verizon and AT&T maintain extensive average 3G coverage levels of 77.1% and 84.1%, respectively. In contrast, T-Mobile and Sprint operate with substantially smaller footprints, averaging just 17.5% and 38.2% coverage across the 145 markets.

Finally, Table 3 breaks down the spatial relationship between pairs of rivals operating under the same technological standard. A clear geometric distinction emerges between the UMTS (AT&T/T-Mobile) and CDMA (Verizon/Sprint) networks. The CDMA pair exhibits a high degree of spatial overlap (averaging 37.22%), indicating that Verizon and Sprint frequently offer coverage in the same geographic blocks. Conversely, the UMTS pair features significantly less direct overlap (17.75%) and a highly asymmetric roaming dynamic.

Table 2
Summary Statistics - Coverage Measures

	Mean	Std. Dev.	Min	25%	Median	75%	%Zeros	N
<i>Coverage Measures</i>								
Total Coverage								
Verizon	77.05	21.83	0.00	66.58	80.39	94.45	3	481
AT&T	84.12	21.92	0.00	76.60	94.58	99.23	1	391
T-Mobile	17.46	19.59	0.00	1.30	10.46	26.99	15	265
Sprint	38.20	29.65	0.00	16.31	29.88	57.62	12	286
Catch-up Area (Rival's Exclusive Coverage)								
Verizon	1.25	3.63	0.00	0.02	0.10	1.36	12	481
AT&T	0.12	0.45	0.00	0.00	0.00	0.02	43	391
T-Mobile	64.84	24.51	0.00	48.12	69.34	84.43	1	265
Sprint	38.30	23.67	0.00	15.26	40.13	53.91	5	286
Own Exclusive Coverage								
Verizon	39.56	22.17	0.00	23.82	40.69	53.42	3	481
AT&T	66.10	23.65	0.00	47.10	70.23	85.93	1	391
T-Mobile	0.11	0.43	0.00	0.00	0.00	0.02	42	265
Sprint	1.44	4.03	0.00	0.01	0.12	1.37	17	286
Competitor Total Coverage								
Verizon	38.74	28.90	0.00	13.70	30.06	57.62	8	481
AT&T	18.14	19.56	0.00	2.36	11.34	28.70	10	391
T-Mobile	82.19	24.62	0.00	75.34	94.24	99.23	1	265
Sprint	75.05	25.00	0.00	64.15	80.39	93.97	5	286
Overlapping Coverage								
Verizon	37.49	28.13	0.00	13.41	29.79	55.38	8	481
AT&T	18.02	19.58	0.00	2.36	10.34	28.70	10	391
T-Mobile	17.35	19.61	0.00	1.30	10.34	26.99	15	265
Sprint	36.75	28.94	0.00	14.63	27.88	55.56	12	286
Empty Area								
Verizon	21.70	21.80	0.00	3.50	17.09	32.80	0	481
AT&T	15.76	21.82	0.00	0.75	5.42	23.36	1	391
T-Mobile	17.70	24.55	0.00	0.75	5.67	23.74	1	265
Sprint	23.51	24.88	0.00	3.58	17.25	34.11	1	286

Note: This table shows summary statistics for the 145 markets over 2015-2018, without accounting for the lagged markets. Coverage variables are in percentage points (0-100).

Table 3
Additional Coverage Measures by Technology Standard

Variable	Mean	Std. Dev.	Min	Median	Max	N
Panel A: UMTS Pair (AT&T / T-Mobile)						
Overlap Coverage (%)	17.75	19.58	0.00	10.34	73.00	656
Empty Area (%)	16.54	22.96	0.00	5.43	100.00	656
Roaming Access (%)	26.26	35.39	0.00	0.05	100.00	656
Rival Roaming (%)	39.44	37.20	0.00	38.03	100.00	656
Panel B: CDMA Pair (Verizon / Sprint)						
Overlap Coverage (%)	37.22	28.42	0.00	28.63	99.13	767
Empty Area (%)	22.38	23.00	0.00	17.09	100.00	767
Roaming Access (%)	15.06	23.20	0.00	1.37	96.58	767
Rival Roaming (%)	25.35	25.58	0.00	15.22	96.58	767

Note: This table shows additional coverage measures computed from pairwise comparisons between firms in the same technology standard. UMTS pair: AT&T and T-Mobile. CDMA pair: Verizon and Sprint. Overlap is the area where both firms have 3G coverage. Empty area is where neither firm has coverage. Roaming access is the area where a firm's customers can roam onto the rival's network.

B Appendix: Empirical Specifications and Additional Results

This appendix provides the full empirical specifications and regression results underlying the stylized facts discussed in Section 3, while also presenting further descriptive evidence.

B.1 Heterogeneity in Coverage and Spatial Investment Incentives




Figure 1

Note: 3G UMTS coverage for the two nationwide providers as of December 2018. Data from FCC Form 477.




Figure 2

Note: 3G CDMA2000 coverage for the two nationwide providers as of December 2018. Data from FCC Form 477.

Firms differ substantially in the geographic allocation of their networks, both across and within technological standards. These differences matter because the incentives to expand coverage depend on the locations that remain unserved and on the coverage provided by compatible rivals. In particular, investment in areas already served by a compatible rival and investment in previously uncovered areas have different implications for competition and roaming.

Figures 1a and 1b display the 3G coverage of the two nationwide providers deploying the UMTS standard, and Figures 2a and 2b display the coverage of the two providers deploying CDMA2000 in 2018. Firms differ markedly in the spatial allocation of their networks both across and within

standards. In each technological path, one provider emerges as a clear coverage leader, while the other exhibits a distinct geographic footprint. Notably, T-Mobile has substantially less coverage than any of the other three nationwide carriers. The extent of this spatial differentiation suggests that the available opportunities for catch-up and empty-area investment vary substantially across markets.

To quantify this heterogeneity, I estimate a set of block-level regressions that capture both the extensive margin, corresponding to the decision to invest at all, and the intensive margin, corresponding to the scale of investment conditional on entry. For each block $c \in m$, I estimate the extensive margin using a Probit model of the form:

$$\Pr(y_{jcm,t} > 0) = \Phi(\beta_1^{\text{ext}} \text{Area}_{jcm,t-1}^{\text{empty}} + \beta_2^{\text{ext}} \text{Area}_{jcm,t-1}^{\text{catch-up}}) \quad (2)$$

and the intensive margin using a Poisson regression conditional on positive investment:

$$\mathbb{E}[y_{jcm,t} \mid y_{jcm,t} > 0] = \exp(\beta_1^{\text{int}} \text{Area}_{jcm,t-1}^{\text{empty}} + \beta_2^{\text{int}} \text{Area}_{jcm,t-1}^{\text{catch-up}}) \quad (3)$$

where $\text{Area}_{jcm,t-1}^{\text{empty}}$ and $\text{Area}_{jcm,t-1}^{\text{catch-up}}$ denote the geographic size of the available unserved and rival-only locations within the block at time $t - 1$, respectively.

Because these regressions are estimated separately for each DMA, they produce a distribution of market-specific marginal effects. Table 4 summarizes these distributions. The baseline returns are generally positive, confirming that larger unserved areas represent greater investment opportunities. However, the catch-up coefficient is systematically larger than the empty-area coefficient across both margins. This difference is consistent with the stronger incentive to invest where a compatible rival already operates, relative to expanding into areas where no compatible carrier is currently present.

The magnitude of this difference is substantial. On the extensive margin, the average coefficient on catch-up area is 0.203, compared with 0.017 for empty area. The corresponding coefficients on the intensive margin are 0.171 and 0.031, respectively. Thus, firms appear to place substantially greater weight on available catch-up opportunities than on empty-area opportunities. This pattern is consistent with the competitive incentive to eliminate a rival's coverage advantage and with the fact that empty-area investment can generate roaming benefits that are partially captured by compatible rivals.

Furthermore, rival scale strongly deters empty area expansion. The penalty derived by compatible networks is directly observable in the cross-section of markets. In markets where a rival carrier holds a larger share of the subscriber base, the focal firm's incentive to invest in empty

areas drops significantly. Because a larger rival captures a greater absolute share of the roaming spillovers, the investing firm’s net return on empty area expansion is severely attenuated. These pronounced cross-market variations confirm that spatial network investments are heavily driven by local competitive forces and compatibility spillovers.

Table 4
Summary of Block-Level Spatial Investment Coefficients

Variable	Mean Coef.	Std. Dev.	% Positive	% Signif. ($p < 0.05$)
<i>Panel A: Extensive Margin (Probit, $N = 129$ DMAs)</i>				
Empty Area (β_1^{ext})	0.0171	0.0400	73%	88%
Catch-up Area (β_2^{ext})	0.2027	0.0445	100%	100%
<i>Panel B: Intensive Margin (Poisson, $N = 137$ DMAs)</i>				
Empty Area (β_1^{int})	0.0306	0.0593	76%	99%
Catch-up Area (β_2^{int})	0.1710	0.0960	100%	100%

Notes: The table summarizes the distribution of DMA-specific coefficient estimates from block-level investment regressions. Panel A reports coefficients from Probit models for the extensive margin, while Panel B reports coefficients from Poisson models estimated conditional on positive investment for the intensive margin. “Empty Area” measures the available unserved area within a block, while “Catch-up Area” measures the area covered by a compatible rival but not by the focal firm at $t - 1$. “% Positive” reports the share of DMA-specific estimates that are positive, and “% Signif.” reports the share that are statistically significant at the 5% level. The number of DMAs differs across margins because the estimation samples differ across specifications.

This systematic difference between empty-area and catch-up investment is further reported in Figures 3 and 4. Figure 3 shows the distribution of the estimated coefficients, highlighting the rightward shift of the catch-up coefficients relative to those for empty areas. Figure 4 plots the empty-area coefficients against the catch-up coefficients for each market, showing that catch-up incentives exceed empty-area incentives across markets.

Figure 5 plots the market-specific point estimates with their 95% confidence intervals for both the extensive and intensive margins, ordered by the magnitude of the empty-area coefficient. The figure shows considerable variation in the estimated incentives to invest across DMAs, with the relative attractiveness of empty-area and catch-up investment differing across local market conditions.

To investigate the sources of this cross-market heterogeneity, I regress the recovered market-specific coefficients ($\hat{\beta}_{1m}, \hat{\beta}_{2m}$) on DMA-level market characteristics. Table 5 presents these second-stage OLS results.

The most interesting result is the strong negative effect of the rival’s market share on the empty-area coefficient (Columns 1 and 3). When the compatible rival has a larger subscriber base, the focal firm faces a steeper positive spillover penalty, as more of the roaming spillovers leak to the competitor. Specifically, a 10 percentage-point increase in the rival’s market share reduces the firm’s empty-area entry coefficient by 0.047 ($p < 0.01$). The relationship provides




Figure 3
Distribution of DMA-Specific Investment Coefficients

Note: The figure reports the distribution of DMA-specific coefficients for available empty-area and catch-up investment opportunities.




Figure 4
Empty-Area and Catch-up Investment Coefficients Across DMAs

Note: Each point represents a DMA-specific pair of estimated coefficients for empty-area and catch-up investment opportunities. The 45-degree line indicates equal coefficients across the two investment margins.

further evidence that the spatial incentives documented above vary systematically with the competitive environment and with the strength of compatibility spillovers.




Figure 5
Within-Market Correlation of Investment Incentives

Note: The figures report DMA-specific coefficient estimates and 95% confidence intervals for the empty-area and catch-up investment margins. Estimates are ordered by the magnitude of the empty-area coefficient.

Table 5
Determinants of Spatial Investment Incentives Across Markets

State Variable	Extensive Margin (Probit)		Intensive Margin (Poisson)	
	Empty Area (β_1^{ext})	Catch-up (β_2^{ext})	Empty Area (β_1^{int})	Catch-up (β_2^{int})
Rival Market Share	-0.4708*** (0.1686)	0.3632* (0.1947)	-0.5210** (0.2377)	0.3995 (0.3608)
Own Market Share	-0.4541** (0.1837)	0.2386 (0.2122)	-0.2801 (0.2438)	1.3030*** (0.3702)
Own 3G Coverage	-0.0568 (0.0401)	-0.0175 (0.0463)	-0.1447*** (0.0499)	0.0153 (0.0758)
Lagged Empty Area	-0.0980*** (0.0349)	0.0151 (0.0403)	-0.1129*** (0.0395)	-0.0372 (0.0599)
Spatial HHI	0.1253* (0.0664)	-0.1100 (0.0767)	0.0587 (0.0961)	-0.1347 (0.1460)
Urban Ratio	-0.0558*** (0.0191)	0.0729*** (0.0220)	-0.1063*** (0.0269)	0.0861** (0.0408)
Price	0.0094*** (0.0034)	-0.0059 (0.0040)	0.0159*** (0.0049)	-0.0191** (0.0074)
Mean Income	0.0298 (0.0210)	-0.0135 (0.0242)	0.0317 (0.0287)	0.0747* (0.0436)
Mean Age	0.0025 (0.0016)	0.0029 (0.0018)	-0.0012 (0.0023)	-0.0003 (0.0035)
Observations	129	129	137	137
Adjusted R^2	0.188	0.126	0.219	0.313

Notes: The table reports OLS regressions of the DMA-specific investment coefficients estimated from the block-level regressions on DMA-level market characteristics. Columns 1–2 use the extensive-margin Probit coefficients as dependent variables, while Columns 3–4 use the intensive-margin Poisson coefficients. Standard errors are reported in parentheses. The number of observations differs across margins because the underlying block-level estimation samples differ across specifications. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.2 Demand and Supply Spillovers

To quantify demand-side spillovers, I study how a focal firm’s market share responds to different types of coverage variations. I distinguish between two channels:

1. **Demand Spillover:** The percentage of the market covered exclusively by the compatible rival, interacted with the rival’s market share in $t - 1$. This captures the supplementary roaming area effectively available to the focal firm’s subscribers.
2. **Business Stealing:** The percentage of the market where only the focal firm has coverage (which enters the rival’s roaming utility), interacted with the rival’s market share in $t - 1$.

To separately quantifying the relationship between each coverage type and the focal firm’s demand, I estimate two separate market-level specifications using the log-odds of the market share of focal firm j relative to the outside option:

$$\ln(s_{jmt}/s_{0mt}) = \alpha + \beta_1 \text{DemandSpillover}_{kj,t-1} + \gamma X'_{jmt} + \delta_j + \theta_t + \kappa_m + \epsilon_{kjt} \quad (4)$$

$$\ln(s_{jmt}/s_{0mt}) = \alpha + \beta_1 \text{BusinessStealing}_{kj,t-1} + \gamma X_{jmt} + \delta_j + \theta_t + \kappa_m + \epsilon_{kjt} \quad (5)$$

Both specifications control for the focal firm’s own coverage, prices, demographics (X_{jmt}), and unobserved heterogeneity through firm (δ_j), year (θ_t), and market (κ_m) fixed effects. In both cases, the coverage measure is scaled by the rival’s lagged market share, reflecting that the magnitude of the roaming effect depends on the number of subscribers who can benefit from the additional coverage.

As shown in Table 6, the results are consistent with the demand-side mechanism described in the main text. The coefficient on the demand spillover is positive and highly significant ($\beta = 8.490$, $p < 0.01$), indicating that greater rival-exclusive coverage is associated with higher demand for the focal firm’s service. In contrast, the coefficient on the compatibility penalty is strongly negative ($\beta = -8.508$, $p < 0.01$). Thus, coverage that expands the effective network available to the focal firm’s subscribers generates a positive demand spillover, whereas coverage that expands the effective network available to the rival generates a negative effect on the focal firm’s demand.

For the supply side, I examine whether the geographic position of the focal firm relative to its compatible rival is associated with different investment responses across catch-up and empty areas. The primary specification interacts the size of the available investment opportunity with the overall **Coverage Gap**, defined as the rival’s exclusive coverage minus the focal firm’s

Table 6
Demand-Side Spillover Estimates (Separate Specifications)

	Model 1	Model 2
	Log-Odds of Market Share	Log-Odds of Market Share
Demand Spillover	8.490*** (0.000)	—
Business Stealing	—	-8.508*** (0.006)
Own Coverage	3.498*** (0.000)	3.597*** (0.000)
Controls	Yes	Yes
Market FE	Yes	Yes
Year FE	Yes	Yes
Firm FE	Yes	Yes
R^2	0.570	0.570

Notes: The table reports OLS estimates from separate market-level regressions of the log-odds of the focal firm’s market share relative to the outside option. Model 1 includes the demand-spillover measure, while Model 2 includes the business-stealing measure. Both specifications control for the focal firm’s own 3G coverage and other market-level covariates, as well as market, year, and firm fixed effects. The demand-spillover measure is defined as compatible-rival-exclusive coverage interacted with the rival’s lagged market share. The business-stealing measure is defined as focal-firm-exclusive coverage interacted with the compatible rival’s lagged market share. P-values are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

exclusive coverage. A positive gap indicates that the rival holds a coverage advantage. The Poisson specification conditional on positive investment is:

$$\mathbb{E}[y_{jmt} \mid y_{jmt} > 0] = \exp(\alpha + \beta_1 \text{Area}_{jmt-1} + \beta_2 (\text{Area}_{jmt-1} \times \text{Gap}_{jmt-1}) + \gamma Z_{jmt} + \mu_j + \tau_t + \nu_{ijt}) \quad (6)$$

where Area_{jmt-1} denotes the relevant available investment area, either catch-up or empty area, depending on the specification.

Table 7 reports the results. The interaction terms indicate that the relative attractiveness of the two investment margins varies systematically with the coverage gap. As the rival’s coverage advantage increases, the marginal response to available catch-up area becomes less negative ($\beta = +0.0001$, $p < 0.01$). At the same time, the marginal response to empty area declines ($\beta = -0.0004$, $p < 0.01$). These patterns are consistent with firms placing greater weight on catch-up opportunities when they face a larger coverage disadvantage, while reducing investment in locations where empty-area expansion generates greater potential roaming benefits for the rival.

Table 7
Intensive Margin Investment Responses to Coverage Gaps (Poisson)

	Model 1 (Catch-up Interaction)	Model 2 (Empty Interaction)
Catch-up Area (lag)	-0.0196*** (0.000)	-0.0188*** (0.000)
Catch-up $\times$ Gap	0.0001*** (0.000)	–
Empty Area (lag)	0.0271*** (0.000)	0.0271*** (0.000)
Empty $\times$ Gap	–	-0.0004*** (0.000)
Coverage Gap	0.0215*** (0.000)	0.0322*** (0.000)
Controls	Yes	Yes
Firm FE	Yes	Yes

Notes: The table reports Poisson estimates for the intensive margin of investment, conditional on positive investment. Model 1 interacts the available catch-up area with the coverage gap, while Model 2 interacts the available empty area with the coverage gap. The coverage gap is defined as the compatible rival's exclusive coverage minus the focal firm's exclusive coverage. Both specifications include the lagged investment opportunity, the coverage gap, market-level controls, and firm fixed effects. P-values are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 6 illustrates these interaction effects on the marginal response of investment. The two margins exhibit sharply different responses to the coverage gap. The marginal response to available catch-up area is initially negative but becomes less negative as the rival's coverage advantage increases (Panel A). In contrast, the positive marginal response to empty-area opportunities declines as the coverage gap increases (Panel B). A one-standard-deviation increase in the coverage gap raises the marginal investment response to catch-up opportunities by 0.309, while reducing the response to empty-area opportunities by 1.297. Thus, when a firm falls further behind its compatible rival, its investment response shifts toward locations where the rival already operates and away from locations where neither firm currently has coverage. This pattern is consistent with the strategic trade-off emphasized in the structural model. Catch-up investment allows a firm to replace roaming access with own coverage and reduce the rival's geographic advantage. Empty-area investment expands the effective network available to the firm's customers but can also create roaming benefits for the compatible rival. The relative attractiveness of these investment margins therefore depends on the firms' existing network positions.

Furthermore, Figure 7 plots the distribution of coverage gaps across the sample alongside the marginal effect of available catch-up area conditional on the coverage gap. The figure provides a visual representation of the relationship between firms' relative network positions and their





Figure 6
Marginal Effects of Catch-up and Empty Areas on Investment by Coverage Gap

Note: The figure plots the marginal effects of available catch-up and empty areas on the intensive margin of investment as a function of the coverage gap. The coverage gap is defined as the compatible rival's exclusive coverage minus the focal firm's exclusive coverage. Shaded areas represent 95% confidence intervals.



(a) Distribution of Coverage Gaps

(b) Base Marginal Effect of Catch-up by Gap

Figure 7
Supplementary Analysis of the Coverage Gap

Note: Panel A shows the distribution of the coverage gap, defined as the compatible rival's exclusive coverage minus the focal firm's exclusive coverage. Panel B plots the marginal effect of available catch-up area on investment as a function of the coverage gap.

investment responses documented in Table 7.

B.3 Overlap and the Spatial Matching Mechanism

When compatible firms expand into previously uncovered areas, their investments may target the same locations. The extent of resulting overlap therefore depends not only on the amount of investment by each firm, but also on the size and characteristics of the area available for expansion. I interpret this process as a spatial matching mechanism and test whether the observed overlap is systematically related to firms' investments and to the characteristics of the locations available for deployment.

To quantify this relationship, I estimate a series of OLS regressions relating overlapping coverage to the investments of focal firm j and its compatible rival k . The specification estimating the *levels* of overlap (Models 1–3) is given by:

$$Overlap_{jk,t} = \alpha + \gamma_1 y_{j,t} + \gamma_2 y_{k,t} + \gamma_3 (y_{j,t} \times y_{k,t}) + \gamma_4 Empty_{jk,t-1} + \theta X'_{jk,t-1} + \varepsilon_{jk,t} \quad (7)$$

where y represents the scale of network investment, *Empty* is the baseline size of the unserved area, and X' captures local demographic controls within the empty area, such as the Spatial HHI and the Urban Ratio.

To capture the formation of new overlap at the margin, I also estimate a *first-differences* specification (Model 4):

$$\Delta Overlap_{jk,t} = \alpha + \gamma_1 \Delta y_{j,t} + \gamma_2 \Delta y_{k,t} + \gamma_3 (\Delta y_{j,t} \times \Delta y_{k,t}) + \nu_{jk,t} \quad (8)$$

Table 8 reports the results. In Model 1, the coefficient on the initial empty area is negative and highly significant ($\gamma_4 = -0.561$, $p < 0.01$), indicating that, conditional on investment, overlap is lower when firms have a larger area over which to deploy their networks. A larger empty area therefore provides more scope for firms to target different locations, reducing the extent to which their investments coincide.

The results also show that local market characteristics systematically shape the spatial allocation of investment. The Urban Ratio has a positive and highly significant association with overlap across the level specifications. In the fully controlled specification (Model 3), the Spatial HHI is also positive and highly significant ($\gamma = 22.001$, $p < 0.01$). These results are consistent with the idea that firms are more likely to target the same locations when profitable deployment opportunities are geographically concentrated. In more concentrated and urban markets, therefore, independent investment decisions are more likely to result in overlapping coverage.

The interaction between the firms' investments provides further evidence on the formation of

Table 8
Determinants of Spatial Overlap (Matching Function Estimates)

Variable	Model 1 (Basic Levels)	Model 2 (+ Concentration)	Model 3 (+ Catch-up Controls)	Model 4 (Δ First-Differences)
Investment i (y_i)	0.311***	0.315***	0.280***	0.060***
Investment j (y_j)	0.262***	0.273***	0.245***	0.053***
Interaction ($y_i \times y_j$)	0.009	0.008	0.038***	-0.049***
Empty Area (Search Friction)	-0.561***	-0.565***	-0.410***	—
Spatial HHI	—	2.236	22.001***	—
Urban Ratio	—	13.585***	11.810***	—
Catch-up Area i	—	—	-1.130***	—
Catch-up Area j	—	—	-1.670***	—
Fixed Effects	Yes	Yes	Yes	Yes
R^2	0.239	0.795	0.954	0.550

Notes: Models 1–3 estimate the level of overlapping coverage. Model 4 estimates year-over-year changes (Δ). Slow-moving demographic characteristics are differenced out in Model 4. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

overlap. In the fully controlled levels specification (Model 3), the positive interaction ($\gamma = 0.038$, $p < 0.01$) indicates that joint investment is associated with more than additive increases in total overlap. In contrast, the interaction is negative in the first-differences specification (Model 4, $\gamma = -0.049$, $p < 0.01$), indicating that the amount of new overlap generated by simultaneous investment increases less than proportionally with the two firms' marginal investments. This pattern is consistent with diminishing marginal overlap formation when both firms are actively expanding into the same set of locations.

Figure 8 provides a graphical summary of the mechanisms estimated in Table 8. It shows how market characteristics, investment, and the resulting overlap are related across markets. The positive relationship between population concentration and overlap is consistent with the interpretation that geographically concentrated investment opportunities increase the likelihood that independent firms target the same locations.




Figure 8
Market Characteristics, Total Investment, and Spatial Overlap

Note: The figure summarizes descriptive relationships between market characteristics, network investment, and overlapping coverage. Panel A plots total investment against population concentration in the pair-specific empty area, with observations differentiated by the Urban Ratio. Panel B reports total investment across terciles of the Urban Ratio in the pair-specific empty area. Panel C plots lagged overlapping coverage against population concentration in the pair-specific empty area, with observations differentiated by total investment. Panel D plots lagged overlapping coverage against total investment separately by quartiles of the lagged pair-specific empty area. Spatial HHI measures population concentration across blocks in the empty area, while the Urban Ratio measures the share of empty-area blocks classified as urban. These variables are computed for each compatible firm pair and DMA-year.

B.4 Overidentification Test of Spatial Nesting (Subset Ratio)

A core premise of the spatial investment framework is the common ranking assumption: because firms face highly correlated local demand shocks, they prioritize expanding into the exact same census blocks. Consequently, a smaller firm's network should geometrically approximate a spatial subset of a larger compatible firm's network.

To formally test this, I compute the *subset ratio* for all compatible pairs, defined as the smaller firm's overlapping coverage with the rival divided by its own total coverage. Under perfect common ranking, this ratio equals 1. The data overwhelmingly support this assumption: across the sample, the mean subset ratio is 0.969, with a median of 0.999. Furthermore, as shown in Figure 9, plotting this ratio against the absolute size asymmetry of the networks ($|y_j - y_k|$) confirms the geometric property. As the size asymmetry between two firms grows, the subset ratio rapidly and tightly converges toward 1.




Figure 9

Overidentification Test: Spatial Subset Ratio vs. Network Asymmetry

Note: The figure reports an overidentification test of the spatial nesting assumption underlying the coverage decomposition. The horizontal axis measures the spatial subset ratio, defined as the share of the focal firm’s coverage that is contained within the compatible rival’s coverage, while the vertical axis measures network asymmetry between the two firms. The relationship is estimated across compatible firm pairs and DMA-years. A systematic relationship between the two measures would indicate that the spatial allocation of coverage is not independent of the relative size and configuration of the two networks.

C Appendix: Micro-Founding the Spatial Matching Mechanism

This appendix provides supplementary evidence justifying the modeling choices made in Section 4.2.2 regarding the first-stage investment and the creation of overlapping areas. Because solving a fully dynamic, block-by-block network entry game is computationally intractable, the structural model approximates spatial collision using a continuous matching function.

Overlap is assumed to arise endogenously from firms independently allocating investment within the available empty area. The efficiency with which these independent investments translate into actual overlapping coverage is governed by a spatial search parameter, $\alpha_{mt,empty}$. Because firms do not deploy base stations randomly, this parameter is modeled as an exponential function of observable local demographics:

$$\alpha_{mt,empty} = \exp(\beta_0 + \beta_1 \text{SpatialHHI}_{mt-1,empty} + \beta_2 \text{UrbanRatio}_{mt-1,empty}) \quad (9)$$

Figure 10 visualizes the cross-sectional identification of this search friction. The strictly downward-sloping relationship between the observed overlap rate and the available empty area illustrates the physical constraint: holding investment scale constant, as the empty area constitutes a larger percentage of the market, the probability of investment collisions systematically drops. The matching parameter α shifts this curve, allowing for markets with concentrated populations to yield higher overlap rates despite having large physical areas.




Figure 10
Identification of the Matching Efficiency Parameter (α)

Notes: The figure illustrates the empirical relationship used to identify the matching efficiency parameter α . The vertical axis reports the overlap rate, defined as the change in overlapping coverage between $t - 1$ and t divided by total network investment in the market-year. The sample is restricted to market-years in which the focal firm invests more than 0.5% of the market area. The horizontal axis measures the share of the pair-specific area that remained unserved in the previous year. Gray points represent individual market-year observations, while red points show means within ten equally spaced bins of the lagged empty-area measure; error bars denote 95% confidence intervals. Variables are normalized to the unit interval. The regression underlying the figure controls for lagged competitor coverage in addition to lagged empty area.



Figure 11
Theoretical Behavior of the Matching Function under varying α

Notes: The figure illustrates the theoretical properties of the spatial matching function for alternative values of the matching efficiency parameter α . The overlap function is given by $O = X_{\text{empty}}[1 - \exp(-y_i/(\alpha X_{\text{empty}}))][1 - \exp(-y_j/(\alpha X_{\text{empty}}))]$, holding the available empty area fixed at $X_{\text{empty}} = 17$. Larger values of α correspond to lower matching efficiency and therefore lower overlap for given investment levels. Contour lines represent combinations of the two firms' investments that generate the same level of overlap.

To further illustrate how this parameter disciplines the structural model, Figure 11 plots the theoretical contours of the matching function. The contours show the expected overlap generated by a given pair of investments (y_j, y_k) . In high- α markets (e.g., highly urban markets with concentrated populations), the contour lines are tighter, meaning that relatively low levels of joint investment rapidly produce spatial overlap. In low- α markets (e.g., markets with diffuse, rural populations), the probability of finding the exact same profitable blocks diminishes, requiring massive scale from both firms to generate the same level of overlap. This continuous formulation perfectly captures the empirical realities of network deployment while preserving tractability for the first-stage estimator.

D Appendix: Proofs of Identification

Proof of Proposition 1.

Proof. Consider a market m and period t with two compatible firms, $\mathcal{J}_r = \{1, 2\}$. Define the effective search effort of firm j as

$$E_{jmt} \equiv \frac{y_{E,jmt}}{\alpha_{o,m} X_{Empty,m,t-1}}, \quad (10)$$

so that

$$y_{E,jmt} = \alpha_{o,m} E_{jmt} X_{Empty,m,t-1}, \quad y_{C,jmt} = y_{jmt}^3 - y_{E,jmt}. \quad (11)$$

The matching function in equation 7 in the paper then implies that expected overlapping coverage is

$$o_{m12t} = X_{Empty,m,t-1} (1 - e^{-E_{1mt}}) (1 - e^{-E_{2mt}}). \quad (12)$$

The observed total investment differs from the equilibrium total anticipated by the firm because the realization of deployment is not perfectly predictable at the time the investment decision is made. In particular,

$$y_{jmt}^3 = y_{E,jmt}^* + y_{C,jmt}^* + e_{jmt}, \quad \mathbb{E}[e_{jmt} \mid \mathcal{I}_{jrt-1}] = 0, \quad (13)$$

where $(y_{E,jmt}^*, y_{C,jmt}^*)$ are the equilibrium allocations and e_{jmt} is the corresponding expectational error. Because the structural objects are measurable with respect to $\mathcal{I}_{jrt-1}$, identification therefore proceeds through the conditional expectation of the observed outcomes.

For the interior allocation, define the latent vector

$$\theta_m = (E_{1mt}, E_{2mt}, \alpha_{o,m}).$$

The observed spatial transition and the two firm-level optimality conditions define the structural system

$$\mathbf{H}(z; \theta_m) = \begin{bmatrix} \Delta X_{\text{overlap}, m12t} - X_{\text{Empty}, m, t-1}(1 - e^{-E_{1mt}})(1 - e^{-E_{2mt}}) \\ MR_{E,1}(\theta_m) - MR_{C,1}(\theta_m) \\ MR_{E,2}(\theta_m) - MR_{C,2}(\theta_m) \end{bmatrix} = \mathbf{0}. \quad (14)$$

Here the marginal revenues are the per-market objects implied by the investment first-order conditions:

$$MR_{E,j}(\theta_m) = (p_{jrt} - mc_{jrmt}) \left[s_{jrmt}(1 - s_{jrmt})\Gamma_{i,3G} - \sum_{k \neq j \in \mathcal{J}_r} \Gamma_{i,R} s_{jrmt} s_{krmt} e^{-E_{kmt}} \right] \frac{e^{-E_{jmt}}}{\alpha_{o,m}}, \quad (15)$$

$$MR_{C,j}(\theta_m) = (p_{jrt} - mc_{jrmt}) s_{jrmt} (1 - s_{jrmt}) (\Gamma_{i,3G} - \Gamma_{i,R}). \quad (16)$$

The empty-area conversion rate is

$$\frac{\partial x_j^{3G}}{\partial y_{E,j}} = \frac{e^{-E_{jmt}}}{\alpha_{o,m}},$$

whereas the catch-up conversion rate is one. The second and third rows of (14) follow from the fact that, at an interior optimum,

$$MR_{E,j} = MC_{jrmt} = MR_{C,j}.$$

Thus, the common marginal cost, including the structural cost shock v_{jrmt} , cancels from the split-equality condition.

Under Assumption 1(iii), the empty-area marginal revenue is strictly positive. To see this, factor s_{jrmt} from the bracket in (15). The remaining term is

$$(1 - s_{jrmt})\Gamma_{i,3G} - \sum_{k \neq j} \Gamma_{i,R} s_{krmt} e^{-E_{kmt}}.$$

Since $e^{-E_{kmt}} \in (0, 1]$ and $\sum_{k \neq j} s_{krmt} \leq 1 - s_{jrmt}$,

$$\sum_{k \neq j} \Gamma_{i,R} s_{krmt} e^{-E_{kmt}} \leq \Gamma_{i,R} \sum_{k \neq j} s_{krmt} \leq \Gamma_{i,R} (1 - s_{jrmt}) < \Gamma_{i,3G} (1 - s_{jrmt}),$$

where the final strict inequality uses $\Gamma_{i,3G} > \Gamma_{i,R}$ from Assumption 1(ii). The bracketed term is therefore bounded below by $(1 - s_{jrmt})\Gamma_{i,3G} - \Gamma_{i,3G}(1 - s_{jrmt}) = 0$ and is strictly positive.

Together with a positive markup, this implies $MR_{E,j} > 0$.

The spatial transitions alone do not identify the three latent components of θ_m . Substituting $y_{C,j} = y_j^3 - y_{E,j}$ into the transition equations gives

$$\Delta X_{\text{unique},1} = y_{E,1} + y_{E,2} - y_2^3 - o_{m12t}, \quad (17)$$

$$\Delta X_{\text{unique},2} = y_{E,1} + y_{E,2} - y_1^3 - o_{m12t}. \quad (18)$$

These unique-coverage transitions are the roaming-coverage transitions of the main text under the identity $\Delta X_{\text{Roam},mj} \equiv \Delta X_{\text{unique},k}$: the coverage that is unique to firm k is precisely the coverage firm j accesses through roaming. I use the two labels interchangeably.

Hence, conditional on observed total investments, the two unique-coverage transitions differ only by the observed constant $y_1^3 - y_2^3$. The three spatial transitions therefore depend on the latent vector only through the overlap o_{m12t} and the sum $y_{E,1} + y_{E,2}$, so that the Jacobian of the physical transition map has rank at most two. The spatial accounting alone consequently cannot separate the two latent allocations from the spatial friction.

The optimality conditions provide the additional restrictions. By equation (13), the conditional expectation of the observed total investment equals its equilibrium value. Hence, at the true parameter θ_{m0} , the conditional expectations of the observed transitions satisfy

$$\mathbf{H}(z; \theta_{m0}) = \mathbf{0}.$$

Assumption 1(i) implies that $\mathbf{H}$ is continuously differentiable in θ_m , while Assumption 1(iv) implies that

$$D_\theta \mathbf{H}(z; \theta_{m0})$$

is nonsingular. The Implicit Function Theorem therefore implies that θ_{m0} is locally isolated and uniquely determined by the conditional expectation of the observables. Thus,

$$(E_{1mt}, E_{2mt}, \alpha_{o,m})$$

is locally point-identified from

$$(\mathbb{E}[\Delta X_{\text{overlap}} \mid \mathcal{I}], \mathbb{E}[y_1^3 \mid \mathcal{I}], \mathbb{E}[y_2^3 \mid \mathcal{I}]).$$

The same argument applies when one of the investment bounds binds. If $y_{E,j} = 0$, the corre-

sponding latent effort satisfies $E_j = 0$; if $y_{C,j} = 0$, then

$$E_j = \frac{y_j^3}{\alpha_{o,m} X_{Empty,m,t-1}}.$$

The corresponding split-equality condition is replaced by the appropriate one-sided Karush–Kuhn–Tucker condition. On each active-set regime, the resulting reduced system is continuously differentiable in the remaining free latent parameters, and the rank condition in Assumption 1(iv) delivers local point identification by the same argument. Assumption 1(ii) rules out simultaneous binding of both bounds, so the relevant regimes are exhaustive and mutually exclusive.

Once $(E_{1mt}, E_{2mt}, \alpha_{o,m})$ is identified, the latent investment allocations follow directly:

$$y_{E,jmt}^* = \alpha_{o,m} E_{jmt} X_{Empty,m,t-1}, \quad y_{C,jmt}^* = y_{jmt}^3 - y_{E,jmt}^*. \quad (19)$$

The spatial friction is parameterized as

$$\alpha_{o,m} = \exp(\mathbf{W}_m^\alpha \beta),$$

where $\mathbf{W}_m^\alpha$ contains the Spatial HHI and Urban Ratio of the empty area. Since $\alpha_{o,m}$ is identified for each market and the covariates $\mathbf{W}_m^\alpha$ have full column rank across markets,

$$\ln \alpha_{o,m} = \mathbf{W}_m^\alpha \beta$$

identifies β .

For completeness, when $\mathcal{J}_r$ contains more than two compatible firms, the same logic extends by combining the pairwise overlap transitions with the firm-level optimality conditions. With J_r compatible firms, there are $J_r + 1$ latent objects,

$$(E_1, \dots, E_{J_r}, \alpha_{o,m}),$$

and the pairwise overlap constraints share the common friction parameter $\alpha_{o,m}$. Each firm's optimality condition contains the corresponding sum of pairwise spillover terms,

$$\sum_{k \neq j} \Gamma_{i,R} s_{jrm} s_{krmt} e^{-E_{kmt}}.$$

The resulting system contains a square subsystem in the $J_r + 1$ latent parameters. The same

Jacobian rank condition delivers local point identification, while the remaining pairwise transitions provide over-identifying restrictions. In the empirical application, each compatible market contains exactly two firms, so the two-firm case above is the relevant one. $\square$

Proof of Proposition 2.

Proof. Proposition 1 identifies the spatial friction and the latent investment allocations. I now show how the level of marginal cost, the cost shock v_{jrm} , and the structural cost parameters λ are recovered.

The firm's two interior first-order conditions are

$$MR_{E,j}(y_E, y_C; \beta) = \mathbf{W}_{jrm} \lambda + v_{jrm}, \quad MR_{C,j}(y_E, y_C; \beta) = \mathbf{W}_{jrm} \lambda + v_{jrm}. \quad (20)$$

Subtracting the two conditions gives

$$MR_{E,j} - MR_{C,j} = 0.$$

Thus the cost level, including v_{jrm} , does not enter the split condition. It is recovered from the level of total investment.

Fix (λ, β) and a candidate value of v_{jrm} . Let

$$J_{\text{FOC}} = \begin{bmatrix} \frac{\partial MR_{E,j}}{\partial y_E} & \frac{\partial MR_{E,j}}{\partial y_C} \\ \frac{\partial MR_{C,j}}{\partial y_E} & \frac{\partial MR_{C,j}}{\partial y_C} \end{bmatrix} = \nabla_{(y_E, y_C)}^2 \pi_j.$$

At an interior optimum, the second-order condition implies that J_{FOC} is negative definite and therefore nonsingular. The Implicit Function Theorem consequently yields a locally unique allocation

$$(\hat{y}_E(v; \theta), \hat{y}_C(v; \theta)), \quad \theta = (\lambda, \beta),$$

that is continuously differentiable in v .

Define the corresponding optimal total investment as

$$\hat{y}^3(v; \theta) = \hat{y}_E(v; \theta) + \hat{y}_C(v; \theta).$$

Because the cost shock enters both first-order conditions with unit coefficient, differentiat-

ing (20) with respect to v gives

$$J_{\text{FOC}} \begin{bmatrix} d\hat{y}_E/dv \\ d\hat{y}_C/dv \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Hence

$$\frac{d\hat{y}^3}{dv} = \mathbf{1}' J_{\text{FOC}}^{-1} \mathbf{1}.$$

Since J_{FOC} is negative definite, its inverse is also negative definite, implying

$$\frac{d\hat{y}^3}{dv} < 0.$$

The firm's optimal total investment is therefore strictly decreasing in the cost shock.

Using the observation equation,

$$y_{jmt}^3 = \hat{y}^3(v_{jrmt}; \theta) + e_{jmt},$$

and the conditional mean-zero property of e_{jmt} gives

$$\mathbb{E}[y_{jmt}^3 \mid \mathcal{I}_{jrt-1}] = \hat{y}^3(v_{jrmt}; \theta).$$

Because $\hat{y}^3(\cdot; \theta)$ is strictly monotone, it is locally invertible, so

$$v_{jrmt}(\theta) = \hat{y}^{3,-1}(\mathbb{E}[y_{jmt}^3 \mid \mathcal{I}_{jrt-1}]; \theta). \quad (21)$$

Thus the cost shock is recovered from the level of equilibrium investment. This inversion is the supply-side counterpart of the demand-side inversion of mean utility from observed market shares.

Finally, substituting the recovered allocation and cost shock into either level first-order condition yields

$$MR_{C,j}(\hat{y}_E(v; \theta), \hat{y}_C(v; \theta)) = \mathbf{W}_{jrmt} \lambda + v_{jrmt}(\theta).$$

Equivalently,

$$v_{jrmt}(\theta) = MR_{C,j}(\hat{y}_E(v; \theta), \hat{y}_C(v; \theta)) - \mathbf{W}_{jrmt} \lambda.$$

Suppose that the cost instruments $\mathbf{Z}_{jrmt}^c$ satisfy

$$\mathbb{E}[\mathbf{Z}_{jrmt}^c v_{jrmt}(\theta_0)] = \mathbf{0},$$

and that

$$\mathbb{E} [\mathbf{Z}_{jrm}^c \mathbf{W}_{jrm}]$$

has full column rank. The resulting moment conditions therefore have a unique solution for λ . Hence the investment cost parameters are point-identified. $\square$

E Appendix: Computational Algorithm

This appendix details the computational implementation of the sequential estimator. The estimation is conducted in two phases that propagate estimated parameters forward. Phase 1 recovers the demand, marginal-cost, and pricing parameters from the second-stage product-market subgame, in three steps: a pseudo-likelihood estimation of the demand heterogeneity parameters, a linear decomposition of the mean utilities, and the inversion of the pricing first-order condition for marginal costs. Phase 2 takes these estimates as given and estimates the investment costs (λ) and spatial-friction parameters (β) from the first-stage network investment game via the MPEC of Section 6.1. Standard errors are constructed to account for the generated-regressors structure across phases.

Phase 1: Second stage estimation

First, I recover the heterogeneity parameters θ_2 governing consumer preferences and the market-provider-year mean utilities δ_{jrm} serving as fixed effects. The algorithm nests an inner-loop share-matching routine inside an outer-loop pseudo-likelihood optimization.

Inner loop: share matching. The standard [Berry \(1994\)](#) approach equates model-predicted shares to observed shares via a logarithmic contraction mapping. However, because the NielsenIQ micro-data relies on survey sampling, the observed weighted market share $\hat{s}_{jrm}$ may be exactly zero for certain carriers in small demographic cells, at which the contraction is undefined. Following [Goolsbee and Petrin \(2004\)](#), I replace the contraction with a Non-Linear Least Squares routine that solves

$$\hat{\delta}_{mt}(\theta_2) = \arg \min_{\delta_{mt}} \sum_{j \in \mathcal{J}_{mt}} (s_{jrm}(\theta_2, \delta_{mt}) - \hat{s}_{jrm})^2,$$

where the predicted share s_{jrm} is the integral of P_{ijrm} over the empirical distribution of consumer demographics. Integrating over the observed consumers, rather than over a parametric demographic distribution, ensures that the mean utilities the inner loop recovers are consistent

with the same individual choices that enter the likelihood.² The [Berry \(1994\)](#) existence and uniqueness result for δ carries over to the NLS formulation, which solves for the same fixed point as the contraction whenever observed shares are strictly positive.

Outer loop: weighted pseudo-likelihood. Taking the $\hat{\delta}_{mt}(\theta_2)$ from the inner loop, the outer loop maximizes the weighted pseudo log-likelihood

$$\max_{\theta_2} \mathcal{L}(\theta_2) = \sum_m \sum_{i \in m} w_i \sum_{j \in \mathcal{J}_{mt}} \mathbf{1}\{y_{ijrmt} = 1\} \ln P_{ijrmt}(\theta_2, \hat{\delta}_{mt}(\theta_2)),$$

where w_i denotes the NielsenIQ projection weights. Optimization uses a quasi-Newton BFGS algorithm with analytic gradients; convergence is assessed by the norm of the gradient and the relative change in $\mathcal{L}$.

I then decompose the mean utilities $\hat{\delta}_{jrmt}$ from Phase 1 into observable product attributes, the national price schedule, and the structural residual ξ_{jrmt} . Because the national price p_{jrt} is endogenous to ξ_{jrmt} , I instrument using 2SLS:

1. *First stage:* Project p_{jrt} on the instrument vector Z_{jrt} and all exogenous covariates and fixed effects to obtain $\hat{p}_{jrt}$.
2. *Second stage:* Regress $\hat{\delta}_{jrmt}$ on $\hat{p}_{jrt}$, coverage X_{jrmt} , and fixed effects $(\omega_m, \tau_t, \Delta_j)$ to recover $\hat{\alpha}$ and $\hat{\gamma}$.
3. *Residual extraction:* Compute $\hat{\xi}_{jrmt} = \hat{\delta}_{jrmt} - \hat{\alpha}\hat{p}_{jrt} - \hat{\gamma}X_{jrmt} - \hat{\tau}_t - \hat{\omega}_m - \hat{\Delta}_j$.

The recovered $\hat{\xi}_{jrmt}$ enters Phase 2 indirectly as a determinant of the predicted market shares $\hat{s}_{jrmt}$, which drive each firm's investment incentives through the marginal-revenue functions.

Given $\hat{\theta}^d = (\hat{\theta}_2, \hat{\alpha}, \hat{\gamma})$, I analytically compute share derivatives $\partial s_{jrmt} / \partial p_{jrt}$ and invert the national pricing first-order condition to recover slope-weighted average marginal costs:

$$\hat{m}c_{jrt}^{(s')} = p_{jrt} + \frac{\sum_m M S_{mt} \hat{s}_{jrmt}}{\sum_m M S_{mt} \partial \hat{s}_{jrmt} / \partial p_{jrt}}.$$

As established in Section 5.1, the national pricing FOC identifies only the slope-weighted average of market-level marginal costs; the supply-stage MPEC uses $\hat{m}c_{jrt}^{(s')}$ as the carrier-year marginal cost entering the margin $(\hat{p}_{jrt} - \hat{m}c_{jrt}^{(s')})$ in every market. This is consistent with the wireless industry's cost structure, in which the dominant operating costs, spectrum licenses, core network

²Following [Goolsbee and Petrin \(2004\)](#) (footnote 27), I recover δ_{mt} by a non-linear least squares share-match rather than by a logit contraction mapping. The two coincide when observed shares are strictly positive, but the least-squares formulation remains well defined when a carrier records a zero NielsenIQ share in a small DMA cell, which the finite NielsenIQ panel occasionally produces.

infrastructure, and corporate operations, are incurred at the national level, while market-specific cost variation is small relative to national margins.

Table 9

Algorithm 1: Phases 1 – Demand and Marginal Costs

Note: The non-linear parameter vector θ_2 contains 10 parameters: three demographics interactions (price $\times$ income, coverage $\times$ age, coverage $\times$ hhsz) and five random coefficient standard deviations (σ).

- 1 **Initialize** heterogeneity parameters $\theta_2^{(0)}$ and mean utilities $\delta_{jrm}^{(0)}$.
 - 2 **Inner loop:** for each θ_2 , solve the share-matching NLS $\hat{\delta}_{mt}(\theta_2) = \arg \min_{\delta_{mt}} \sum_j (s_{jrm}(\theta_2, \delta_{mt}) - \hat{s}_{jrm})^2$ (Goolsbee and Petrin, 2004).
 - 3 **Outer loop:** maximize the weighted pseudo log-likelihood $\mathcal{L}(\theta_2)$ via a quasi-Newton (BFGS) optimizer, with $\hat{\delta}_{mt}(\theta_2)$ supplied at each evaluation.
 - 4 **Decompose** $\hat{\delta}_{jrm}$ via 2SLS into linear parameters $(\hat{\alpha}, \hat{\gamma})$ and the residual $\hat{\xi}_{jrm}$.
 - 5 **Invert** the national pricing FOC to recover slope-weighted marginal costs $\hat{m}c_{jrt}^{(s')}$.
 - 6 **Compute** the marginal-revenue components and pass to Algorithm 2.
-

Phase 2: First-Stage MPEC Estimation

The supply-side estimation recovers the structural parameters (λ, β) via Generalized Method of Moments (GMM), minimizing a criterion that combines cost-side moments, overlap-equation moments, and an inversion penalty. This is implemented via Mathematical Programming with Equilibrium Constraints (MPEC), following Su and Judd (2012) and Dubé et al. (2012). Rather than solving for the equilibrium network allocation in a nested inner loop, the MPEC treats the structural parameters, the latent equilibrium objects $(\mathbf{y}_E, \mathbf{y}_C, \mathbf{o})$, and the cost shocks $\mathbf{v}$ as simultaneous decision variables for the optimizer. Proposition 1 guarantees that the equilibrium mapping is locally unique, ensuring the MPEC computes the same estimator as a nested fixed-point algorithm (identical up to solver tolerance), but without the computationally prohibitive inner loops.

The estimator searches simultaneously for λ and β such that the entire system, the first-order condition (FOC), the overlap equations, and the constructed equilibrium investments, reproduces both the observed total investments and the moment orthogonality conditions.

At any candidate parameter guess (λ, β) , the latent area allocations $(\hat{y}_E, \hat{y}_C)$ are not free parameters. Economically, these allocations are pinned by the FOCs; the MPEC imposes that pinning as a constraint rather than solving it in an inner loop. They are constructed by requiring that the firm’s marginal returns equalize across both margins and equal the marginal cost:

$$MR_{E,jmt}(\hat{y}_E; \beta, \text{rival}) = W_\lambda \lambda + v_{jrm}, \quad (22)$$

$$MR_{C,jmt}(\hat{y}_C) = W_\lambda \lambda + v_{jrm}. \quad (23)$$

While both margins share the same marginal cost, their marginal returns operate differently. The catch-up margin MR_C is straightforward: catch-up area is, by definition, territory where the rival already provides coverage. Investing there simply upgrades the firm's own roaming users to native coverage, with no rival spillover.

The empty margin MR_E is not trivial: new empty-area coverage may or may not overlap with the rival's empty-area deployment, and where it does not, it grants the rival a free roaming upgrade. This spillover is weighted by the overlap probability, which itself depends on the search efficiency $\alpha_{o,m}(\beta)$. Constructing the empty-area return therefore requires solving the FOC jointly for both the allocation and the spatial-friction parameters β . The algorithm searches for the allocation $(\hat{y}_E, \hat{y}_C)$ and the parameters β that jointly satisfy the first-order conditions.

Because the map expansions (ΔX) are fixed by the data, any candidate guess for $\alpha_{o,m}(\beta)$ dictates the exact combination of $\hat{y}_E$ and $\hat{y}_C$ that can satisfy the physical accounting of the map. The constraints admit a unique allocation consistent with the map accounting, at which the joint FOC system holds (exactly at interior allocations, and as a one-sided condition where a bound binds).

The structural cost shock v_{jmt} is the residual making both FOCs hold jointly at the predicted allocation $(\hat{y}_E, \hat{y}_C)$. The sum of these predicted allocations need not match the observed total investment y_{jmt}^3 exactly: the two differ by a forecast error e_{jmt} , unobserved by both the firm and the econometrician. The inversion anchors the constructed total to the observed data through the relation $\hat{y}_E + \hat{y}_C + e_{jmt} = y_{jmt}^3$, with e_{jmt} recovered as the residual of the minimum-distance fit.

To understand the mechanics, consider how a standard nested algorithm with an inner and outer loop would operate. The inner loop computes the investment shocks v . In doing so, it first recovers the predicted total investment by solving the Nash FOC system for $(\hat{y}_E, \hat{y}_C)$ at any candidate (λ, β) . Because the marginal return to empty-area investment depends on the probability of overlapping with the compatible rival, the optimal split between empty and catch-up area must simultaneously satisfy the state transitions and the overlap equation for any candidate β . The shock v is then computed as the difference between observed and predicted total investment, the latter defined as the sum $\hat{y}_E + \hat{y}_C$; at this stage the routine also forms the residual from the search-friction equation 8 in the paper. In the outer loop, I construct the two sets of moments, one for investment cost, one for the overlap equation, evaluate the GMM objective $g'Wg$, and update (λ, β) by minimum distance. The forecast error e is the difference between the observed total investment and the firm's true investment.


The MPEC approach completely flattens this architecture. It promotes the allocations $(\hat{y}_E, \hat{y}_C)$,

the overlaps o , and the cost shocks v to decision variables inside the optimizer alongside (λ, β) , and imposes the equilibrium conditions, state transitions, and overlap equation as constraints. Subject to these constraints, the solver minimizes the GMM objective $g(v)'Wg(v)$. At the joint solution, the overlap matching residual c^B —the gap between observed and constructed $\Delta X_{\text{overlap}}$ —disciplines $\alpha_{o,m}$, tracing out and identifying β , while λ is identified by the orthogonality of the recovered cost shocks v to the exogenous instruments $\mathbf{Z}^c$.

Table 10

Phase 2 – MPEC Estimation of the Investment Game

A single constrained problem: the optimizer chooses (λ, β) and the latents (y_E, y_C, o, v, e) simultaneously. Blocks (i)–(iii) are the constraints; block (iv) is the objective minimized subject to them. Arrows denote logical dependence, not an order of execution.



Standard Errors

I report standard errors for the demand parameters $\hat{\theta}^d$ and the investment-stage parameters $(\hat{\lambda}, \hat{\beta})$. The estimator is sequential: the demand parameters, together with the implied margins, shares, and elasticities, are estimated in Phase 1, and the investment-stage parameters are

estimated by the MPEC of Section 6.1 conditional on these first-stage inputs. I compute standard errors for each stage in turn.

Demand-stage standard errors. Standard errors for the demand parameters $\hat{\theta}^d = (\hat{\theta}_2, \hat{\alpha}, \hat{\gamma})$ are computed following Goolsbee and Petrin (2004), who derive the sandwich covariance matrix for two-stage random-coefficient logit estimators using micro-data with survey-based sampling. This procedure accounts for the NielsenIQ sampling design and propagates the uncertainty from the Phase 1 pseudo-likelihood into the 2SLS decomposition, delivering standard errors for the full demand parameter vector.

Supply-stage standard errors. Standard errors for the investment-stage parameters $(\hat{\lambda}, \hat{\beta})$ are computed by a pair-level block bootstrap, conditional on the first-stage demand estimates. The compatible pair $(j, k) \in \mathcal{J}_r$ is the natural resampling block: the matching moment is defined on pairs and the state-transition constraints link the two firms' allocations within a pair, so drawing whole pairs preserves the within-pair dependence in overlaps and transitions that the estimator exploits, whereas resampling individual firm–market observations would break these relationships and distort the matching moment. In each replication I draw compatible pairs at random and with replacement until I obtain the same number of pairs as in the estimation sample; the pair-level matching moment is formed over the drawn pairs, including repeats, and the firm-level cost moment over the distinct firms belonging to those pairs. Holding the demand estimates $\hat{\theta}^d$ and the implied margins, shares, and elasticities fixed at their point values, I rebuild the corresponding data and index structures and re-solve the MPEC, warm-started at the point estimates, to obtain $(\hat{\lambda}_b, \hat{\beta}_b)$. I use $B = 2000$ replications. Standard errors are computed from the bootstrap distribution, and I report the associated percentile 95% confidence intervals $[Q_{2.5}, Q_{97.5}]$.

Because the supply bootstrap holds the first-stage demand estimates fixed, the resulting standard errors are conditional on the demand stage and do not incorporate first-stage sampling uncertainty; they should therefore be read as a lower bound on the unconditional standard errors of $(\hat{\lambda}, \hat{\beta})$.³

³Propagating first-stage uncertainty into the investment-stage parameters would require either re-estimating the demand stage within each bootstrap replication or a Murphy–Topel-type correction that carries the demand covariance through the supply moments; I do not pursue this here, as re-estimating the random-coefficient demand model within each replication is computationally prohibitive.

F Appendix: Additional Results and Robustness

This appendix presents supplementary results and diagnostics for the two-stage investment model. It reports the full set of estimates underlying the main-text results, including marginal costs, markups, price and coverage elasticities, fixed effects, investment allocation, spatial matching, and external validation of the estimated investment costs.

F.1 Marginal Costs, Markups, and Price Elasticities

Table 11 presents a summary of the profit margins recovered from the estimates of the preferred model reported in Table 2 of the paper. The estimated marginal costs capture the cost of serving an additional subscriber and are distinct from the costs of network deployment modeled in the first stage. This distinction is important for the counterfactual analysis, where changes in compatibility affect both consumers’ effective network quality and firms’ incentives to expand coverage. The estimated markups summarize the profitability of the observed pricing decisions, while the price elasticities characterize the strength of substitution across providers that enters the firms’ pricing first-order conditions.

Estimated service marginal costs range from \$1.68 (AT&T) to \$1.97 (T-Mobile) in data units, corresponding to monthly per-subscriber costs of approximately \$17–\$20. Average markups are 47.4% of price. AT&T and Verizon carry the largest absolute markups (roughly \$1.95/month in data units), consistent with their premium-brand positioning; Sprint and T-Mobile, the market’s smaller carriers, have lower markups and operate closer to marginal cost. Own-price elasticities range from -1.9 (AT&T) to -2.7 (Sprint).

Table 11
Marginal Costs, Markups, and Price Elasticities

Firm	Price ^a	MC ^a	Markup (\$)	Markup (%)	Own-price elasticity
AT&T Mobility	3.62	1.68	1.94	53.6	-1.92
Sprint	3.20	1.86	1.34	41.9	-2.68
T-Mobile	3.35	1.97	1.38	41.2	-2.52
Verizon Wireless	3.76	1.81	1.95	51.9	-2.02
Sample average	3.48	1.83	1.65	47.4	-2.28

Notes: ^a In data units (\$10/month). Multiply by 10 for monthly ARPU and by 120 for annual ARPU. Example: a data-unit price of 3.62 corresponds to a \$36.20/month or \$434/year annual average revenue per user. Step 5 of Phase 1, reported in Table 9, inverts the slope-weighted national Bertrand pricing first-order condition using the demand parameters. Markups are $p_j - \widehat{mc}_j$; markup share is the markup as a percentage of price. Own-price elasticities ε_{jj}^p are computed at the equilibrium shares and prices, averaged across DMAs and years.

F.2 Firm and Time Fixed Effects

Table 12 summarizes the brand (firm) and period (year) fixed effects recovered alongside the linear utility block in Phase 1 using the preferred specification framework. The firm fixed effects favor the CDMA2000 carriers: Verizon and Sprint have higher baseline utilities than the UMTS carriers, with Verizon’s fixed effect the largest.

A plausible explanation relates to the period of analysis and the technological characteristics of the standards. Over these years consumers shifted sharply from voice toward data usage,⁴ and because the CDMA2000 deployments supported better data performance over this window, it is reasonable that consumers on average attached higher baseline utility to the CDMA carriers.

Table 12
Linear Phase Fixed Effects and Baseline Diagnostic Tests

Component	Parameter / Reference Group	Coef.	Std. Err.
<i>Firm Brand Fixed Effects (Reference = AT&T Mobility)</i>			
Δ_{Sprint}	Sprint Brand Effect	0.191	0.92
$\Delta_{\text{T-Mobile}}$	T-Mobile Brand Effect	−4.31	1.39
Δ_{Verizon}	Verizon Wireless Brand Effect	2.81	1.00
<i>Temporal Fixed Effects (Reference = Year 2015)</i>			
τ_{2016}	2016 Dummy	1.83	1.65
τ_{2017}	2017 Dummy	−2.89	1.91
τ_{2018}	2018 Dummy	−13.62	2.56
<i>First-Stage Instrumental Variable Diagnostics (Price on Excluded IVs)</i>			<i>F-Statistic</i>
Variant V1: Base Instruments + Auction Spend (t)			7,843.35
Variant V2: Base Instruments + Auction Spend ($t - 1$)			7,541.70
Variant V3: Base Instruments + Auction Spend ($t - 2$)			7,383.82
Variant V4: Base Instruments + Cumulative Auction Spend (Preferred)			12,953.08
Variant V5: Base Instruments + Auction Spend ($t, t - 1$)			6,470.20
Variant V6: Base Instruments Only (No Auction Data)			10,938.06

Notes: Firm and year fixed effects are estimated alongside the linear utility block in Phase 1 of Specification A; the tabulated coefficients are averages across markets and standard errors are clustered at the carrier–year level. The lower panel reports the observation-level (firm–DMA–year) first-stage F for alternative excluded sets. The first-stage F computed on the identifying carrier–year variation for the preferred (V4) set is 8.52. Market fixed effects (144 parameters partialled out) are omitted for brevity.

F.3 Coverage Elasticities and Cross-Effects

I now discuss in detail the coverage elasticities reported in Table 3 in the main text. The estimated demand system also allows me to quantify how consumers respond to changes in network coverage. These coverage elasticities provide a direct link between the demand estimates and the investment incentives modeled in the first stage. Additional own coverage increases the attractiveness of a provider, while the effect of a rival’s coverage depends on whether the rival

⁴As reported in the FCC reports, data usage per smartphone subscriber rose to an average of 5.1 GB per month, an increase of roughly 31% from 2016 to 2017, accompanied by a drop of about 21% in total annual voice minutes and 9% in messaging traffic.

operates under the same technological standard.

For compatible rivals, the cross-effect combines a positive roaming component with a conventional competitive effect. When the rival provides coverage in locations not served by the focal firm, consumers can benefit from that coverage through roaming, increasing the effective value of the focal firm's service. By contrast, when the rival expands coverage in locations already served by the focal firm, the resulting increase in the rival's effective network strengthens its ability to attract consumers, generating a business-stealing effect.

To quantify the economic importance of network quality, I compute the percentage change in a provider's market share resulting from a 1% increase in its 3G geographic coverage. Table 3 in Section 7.1 reports the resulting mean elasticities. The results reveal substantial heterogeneity across firms. AT&T, Verizon, and Sprint exhibit large own-coverage elasticities (6.7, 6.1, and 6.3, respectively), indicating that share responds strongly to coverage improvements for these carriers. T-Mobile displays a markedly lower own-coverage elasticity (2.1), suggesting that its subscribers' choices are comparatively less sensitive to incremental coverage, consistent with its lower baseline utility and its concentration in dense, already-served urban markets where the marginal coverage block adds little reach.⁵

The cross-coverage elasticities are negative throughout, indicating that, in equilibrium, a carrier's coverage expansion draws share away from its rivals. The roaming spillover nevertheless generates an important asymmetry. Within-standard cross-coverage elasticities are smaller in magnitude on average, whereas cross-standard elasticities are larger, with AT&T and Verizon, the two largest carriers, competing most directly. When a compatible rival expands, the investing firm's subscribers can benefit from the new infrastructure through roaming, partially offsetting the competitive loss. When an incompatible rival expands, no such offset exists, leaving the conventional competitive effect as the dominant channel.

The difference is particularly relevant across the two standards. Network spillovers are stronger within the UMTS pair than within the CDMA pair, where Sprint's share remains relatively sensitive to Verizon's coverage. Differences in device- and network-level compatibility provide an important institutional context for this pattern. For 3G voice, AT&T roamed on T-Mobile and Sprint roamed on Verizon, but consumer device portability was asymmetric: unlocked GSM/UMTS handsets could typically move between AT&T and T-Mobile, whereas CDMA devices without SIM cards could not be transferred between Verizon and Sprint. This constraint limited the effective roaming benefit within the CDMA segment, especially early in the sample, and is consistent with the weaker within-standard offset observed for the CDMA carriers relative

⁵Sprint's coverage distribution across the US is reported in Figure 2, Appendix B.

to the UMTS carriers.

These institutional differences should be interpreted as context rather than as a separately estimated mechanism. Under the GSM/UMTS technology used by AT&T and T-Mobile, SIM-based authentication made switching devices and carriers more straightforward, whereas the handset architecture of CDMA2000 networks historically made switching between operators less flexible. The structural model does not separately estimate device portability as an independent choice margin; its demand-side implication is captured through the roaming component of consumer utility. Device portability therefore helps explain why the effective benefits of compatibility and roaming may differ across standards, but it is not itself an additional estimated channel.

F.4 Investment Allocation Across Geographic Margins

Table 13 reports the model-predicted allocation of investment between empty and catch-up areas. On average, firms invest more in previously uncovered areas than in catch-up areas, suggesting that expansion into empty areas remained an important component of network investment during the period under study. At the same time, the allocation reflects the strategic trade-off at the heart of the investment model: catch-up investment improves a firm's native coverage where a compatible rival already operates, whereas empty-area investment expands the firm's footprint but can also generate a roaming benefit for a compatible rival.

The allocation differs across standards. CDMA providers invest substantially more in empty areas than UMTS providers, with mean empty-area investment of 0.0409 compared with 0.0180. In contrast, predicted catch-up investment is similar across standards, at 0.0121 for CDMA and 0.0135 for UMTS. Thus, the difference in total investment across standards is driven primarily by the allocation toward previously uncovered areas rather than by systematic differences in catch-up investment.

This pattern is consistent with the different degree of effective compatibility across the two technologies. Because compatible coverage provides a partial substitute for own deployment, the return to expanding into an area already served by a compatible rival differs from the return to entering an entirely unserved area. The estimated allocation therefore illustrates why the first-stage investment decision cannot be summarized by a scalar level of network quality.

Table 13
Model-Predicted Investment Allocation

Quantity	Mean	Std. Dev.
<i>All observations</i>		
y_{empty}	0.0295	0.0956
y_{catch}	0.0128	0.0360
y_{total} (predicted)	0.0423	0.1033
y_{total} (observed)	0.0419	0.1037
<i>By standard</i>		
y_{empty} UMTS	0.0180	0.0916
y_{empty} CDMA	0.0409	0.0982
y_{catch} UMTS	0.0135	0.0401
y_{catch} CDMA	0.0121	0.0315
y_{total} UMTS	0.0315	0.0994
y_{total} CDMA	0.0530	0.1061

Note: Coverage measured in percentage points of geographic area (0–1 scale). Reported solution.

F.5 Matching Frictions and Spatial Overlap

The first-stage model also estimates the spatial matching process that maps aggregate investment into the geographic allocation of coverage. The matching-friction parameters determine the probability that investments by compatible firms overlap when they target previously uncovered areas.

The estimates imply that overlap is more likely when profitable investment opportunities are spatially concentrated and when the remaining uncovered areas are more urban. The Spatial HHI and Urban Ratio enter the matching function with negative coefficients, implying a lower $\alpha_{o,m}$ and therefore greater overlap in markets where unserved opportunities are concentrated in a small number of locations or where the remaining uncovered areas are predominantly urban. This pattern is consistent with firms independently targeting similar high-value locations when profitable investment opportunities are geographically concentrated.

This mechanism is supported by the reduced-form evidence in Appendix B. Investment by either compatible firm significantly increases subsequent overlap, while a larger initial empty area reduces overlap, consistent with geographic search frictions.

The estimates are also consistent with the spatial evolution of 3G deployment during the period under study. By this stage of deployment, many densely populated areas had already been covered, leaving a larger share of remaining investment opportunities in rural and sparsely populated locations. This geographic dispersion reduces the likelihood that independent investments target exactly the same locations, generating lower overlap even when firms continue to expand their networks.

The estimated concavity of the matching process is particularly relevant for the counterfactual

Table 14
Estimated Matching Friction Parameters

<i>Panel A: Matching Friction Function</i>				
Parameter	Estimate	Std. Err.		
Constant (β_0)	2.215	0.083		
Spatial HHI (β_{HHI})	-2.542	0.182		
Urban Ratio (β_{urban})	-2.134	0.409		
<i>Panel B: Implied Market-Level Matching Friction</i>				
	Mean	Std. Dev.	10th Perc.	90th Perc.
Matching Friction ($\hat{\alpha}$)	7.451	1.647	5.206	9.139

Notes: Panel A reports the estimated coefficients of the matching-friction function. The coefficients β determine the market-specific matching-friction parameter $\alpha_{o,m}$ through Equation 8 in the paper. Panel B reports the distribution of the implied market-level matching-friction parameter $\hat{\alpha}$ across observations.

analysis. Because overlap does not increase proportionally with total investment, a change in aggregate deployment can alter both the amount of coverage and its spatial configuration. The matching function therefore allows the model to distinguish the welfare consequences of investment that creates new exclusive coverage from those of investment that generates overlapping coverage and roaming spillovers.

F.6 External Validation of Investment Costs

The estimated coverage costs can also be compared with external engineering benchmarks for wireless network deployment. This comparison provides an independent reference point for the magnitude of the structural investment costs, which capture the economic cost of increasing geographic coverage rather than the engineering cost of constructing a particular type of network site.

To make the comparison, I translate engineering estimates of the cost of a new base station into an implied cost per percentage point of geographic coverage using the observed relationship between network deployment and coverage expansion in the data. The conversion requires assumptions concerning site density, the geographic reach of a newly deployed base station, equipment costs, and the distinction between new construction and upgrades to existing infrastructure. The resulting benchmark should therefore be interpreted as an external validation exercise rather than as an alternative estimate of the structural investment cost. The detailed engineering sources, assumptions, and calculations are reported in Appendix H.

Under the assumptions described in Appendix H, I find that building a new base station leads to an increase of approximately 17.5% in the covered area within a market.⁶ Combining this coverage increment with engineering estimates of the total cost of a macro base station (Johansson et al., 2004; Claussen et al., 2008; Björkegren, 2019), I obtain an implied cost of approximately

⁶For a detailed explanation of how I compute the base-station-related coverage measure, see Appendix H.

\$16,000–\$36,000 for a one-percentage-point increase in coverage.⁷

The structural estimates range from one-half of the lower bound to a little less than the upper bound of this engineering benchmark. This difference is consistent with the fact that the two measures capture different margins of network investment. The engineering estimates refer to the construction of a new macro site, whereas the structural parameter captures the marginal cost of increasing geographic coverage by one percentage point, regardless of how that additional coverage is generated. Coverage expansion need not require the construction of a new tower: carriers can expand or improve coverage by upgrading and densifying existing sites, including through co-location, sector and antenna upgrades, and spectrum refarming. Such investments can generate additional coverage at a lower cost than a new macro-site build.⁸

The engineering benchmark should therefore be interpreted as a useful external reference rather than a direct estimate of the structural parameter. Moreover, the engineering calculation is conservative in the opposite direction: the coverage gain from a new site is based on a flat-terrain propagation assumption and assumes no decreasing marginal coverage from additional sites. These assumptions tend to increase the estimated coverage gain per site and therefore lower the implied engineering cost per percentage point. The benchmark may thus understate the cost of generating an additional percentage point of coverage through new-site construction, making the comparison with the structural estimates conservative.

F.7 Alternative Specifications of Demand and Instrumental Variables

To assess the robustness of the main estimates and validate the choice of instruments, I report demand estimates using three alternative instrument sets. The baseline specification, Specification A, instruments for price using auction expenditures, the timing of the 2G network deactivation, and the aggregate share of 4G device penetration. Specification B augments this set by adding the timing of each DMA’s initial 4G rollout as additional instruments. Specification C includes lagged own 2G coverage as a demand covariate in the mean-utility equation rather than relying exclusively on the lagged-coverage instrument.

Table 15 reports the results for the three specifications. Specification B reintroduces the 4G rollout instruments (4G LTE population coverage and years until the activation of the 4G network) and fails the Hansen over-identification test ($J = 13.72$, $p = 0.028$). The combined

⁷Claussen et al. (2008) estimate a cost of \$285,000 over a 15-year horizon, while Johansson et al. (2004) report a cost of \$266,000 for a macro base station.

⁸Industry evidence documents the importance of upgrading existing sites alongside new construction. For example, AT&T’s 2010 network investment plans included both new cell sites and substantial upgrades of existing sites to 3G. See AT&T, “AT&T Upgrades 3G Technology at Cell Sites Across Nation,” January 5, 2010, and AT&T, “AT&T Investment in California Network Aimed at Enhancing Mobile Broadband Service Across the State,” February 16, 2010.

Table 15
Specification Comparison

	Spec A (preferred)	Spec B (+ 4G IVs)	Spec C (+ 2G covariate)
Price (SE)	-1.703 (2.473)	-1.347 (2.262)	-1.666 (2.301)
3G Total Coverage	23.493	20.724	25.803
3G Roaming Coverage	2.192	5.732	3.522
4G Coverage	7.188	8.742	10.586
2G Coverage	–	–	-4.437
Log-likelihood (Phase 1)	-1,414.67	-1,424.26	-1,414.67
First-stage F (obs-level)	12,953.1	9,146.8	12,985.0
R^2 (Phase 2)	0.524	0.510	0.525
Average marginal cost (\$)	1.83	1.51	1.79
Average markup (\$)	1.65	1.97	1.69
Average price elasticity	-2.28	-2.12	-2.23
Over-id test J (p)	13.45 (0.001)	13.72 (0.028)	–

Note: All three specifications recover the same demand and cost objects; they differ in the excluded-instrument set and covariates. Spec A is retained as the baseline. Both the combined Spec A instrument set and the expanded Spec B set reject their over-identifying restrictions (Spec A: Sargan $J = 13.45$ on 2 degrees of freedom, $p = 0.001$; Spec B: Hansen $J = 13.72$ on 4 degrees of freedom, $p = 0.028$). These rejections reflect disagreement among the excluded instruments about the level of the mean price coefficient $\bar{\alpha}$, which is weakly identified from the limited between-year variation in the national price, and not instability in the objects that drive the counterfactuals. The observation-level first-stage F 's overstate instrument strength because the national price is replicated across DMAs; the carrier-year first-stage F for Spec A is 8.52. The coverage valuations, marginal costs, markups, and own-price elasticities are stable across all three specifications.

Specification A instrument set also rejects its over-identifying restrictions ($J = 13.45$, $p = 0.001$). Both rejections signal that the mean price coefficient is only weakly identified from the between-year variation in the national price; the coverage, cost, markup, and elasticity estimates nonetheless remain close across the two specifications.

Specification C includes 2G coverage as an additional demand covariate as a robustness check. Because the years remaining until the 2G network deactivation are used as an instrument for both prices and investment costs, the exclusion restriction would be violated if residual 2G coverage directly affected consumer demand in the second-stage demand model. To assess this possibility, I include 2G coverage directly in the mean-utility specification. If the exclusion restriction is valid, the coefficient on 2G coverage should be statistically insignificant and the estimated price coefficient should remain stable. As reported in Table 15, the coefficient on 2G coverage is economically small and the estimated price coefficient remains largely unchanged, providing evidence in support of the exclusion restriction.

Overall, the estimates are remarkably stable across the three specifications. The negative price coefficient, the positive valuation of network coverage, and the recovered marginal costs, which range between \$1.51 and \$1.83 on average as reported in Table 16, vary little across specifications. What moves across instrument sets is the level of the mean price coefficient. Specifica-

tion A is retained as the baseline because it is the most parsimonious instrument set that delivers stable coverage, cost, and elasticity estimates together with a strong first stage on the identifying carrier-year variation; the price sensitivity that governs substitution is micro-identified from the demographic interactions rather than from the level of $\bar{\alpha}$. Moreover, as reported in Table 12, Specification A is the one providing the highest first stage overall F-statistics, compared to other specifications. These estimates are therefore used for all subsequent welfare and counterfactual analyses.

Table 16
Marginal Cost Inversion Across Specifications

Firm	Specification B				Specification C			
	Price	MC	Markup	Elasticity	Price	MC	Markup	Elasticity
AT&T Mobility	3.62	1.12	2.51	-1.54	3.62	1.63	1.99	-1.88
Sprint	3.20	1.83	1.37	-2.73	3.20	1.84	1.37	-2.62
T-Mobile	3.35	2.10	1.25	-2.75	3.35	1.94	1.41	-2.47
Verizon Wireless	3.76	0.99	2.77	-1.45	3.76	1.77	1.99	-1.97
Average MC			1.51				1.79	
Average Markup			1.97				1.69	
Markup (% Price)			56.7%				48.5%	
Average Elasticity			-2.12				-2.23	

Note: The table reports the implied marginal costs obtained from the slope-weighted marginal cost inversion in Equation 14. Specification B corresponds to the demand model estimated with the expanded set of 4G rollout instruments. Specification C includes 2G coverage as an additional demand covariate. Across specifications, the recovered marginal costs, markups, and elasticities remain broadly similar, supporting the robustness of the pricing and cost estimates to alternative demand specifications.

F.8 First Stage Robustness

To assess the sensitivity of the first-stage estimates to the optimization routine, Table 17 reports the results obtained from five different starting values. The estimates converge to similar, although not identical, parameter vectors. In particular, the cost parameters retain the same qualitative ordering across all solutions: UMTS deployment remains less costly than CDMA deployment, and urban coverage expansion remains more expensive than rural expansion for both standards. Importantly, the objective values remain close across the preferred solutions, suggesting that the model consistently converges to economically similar optima rather than fundamentally different explanations of the data.

Panel B of Table 17 shows that this parameter variation translates into only modest differences in the implied investment costs of coverage expansion. The mean costs recovered across the five solutions lie well within the corresponding minimum and maximum values, and the range of variation is relatively small. For example, the implied cost of expanding UMTS coverage in urban areas ranges from 0.95 to 1.02, while CDMA urban costs range from 2.28 to 2.45. Across all starting values, CDMA investments remain substantially more expensive than UMTS

Table 17

Implied Marginal Cost of a One-Percentage-Point Coverage Expansion

Lambda vector	UMTS rural	UMTS urban	CDMA rural	CDMA urban
<i>Panel A: implied marginal cost by multistart solution</i>				
Start 1 (= selected $\hat{\lambda}$)	0.5905	0.9745	1.8208	2.4402
Start 2	0.5385	1.0137	1.8188	2.3839
Start 3	0.5365	0.9875	1.8799	2.4512
Start 4	0.4834	0.9549	1.8507	2.2839
Start 5	0.6253	1.0204	1.8876	2.3996
<i>Panel B: summary across the five starts</i>				
Mean	0.5549	0.9902	1.8516	2.3918
Min	0.4834	0.9549	1.8188	2.2839
Max	0.6253	1.0204	1.8876	2.4512

Notes: λ -only component of the marginal cost of one percentage point of coverage, in model cost units; the full marginal cost adds the firm–market–year cost shock. Start 1 is the selected estimate used throughout the analysis; the remaining columns are alternative multistart solutions, reported for robustness.

investments, and urban deployment remains more costly than rural deployment. Thus, the main economic conclusions of the first stage are robust to the choice of starting values.

Table 18

Per-Market Cost of a One-Percentage-Point Coverage Expansion

Lambda vector	UMTS rural	UMTS urban	CDMA rural	CDMA urban
Start 1 (= selected $\hat{\lambda}$)	8,188.2	13,511.8	25,247.1	33,834.7
Start 2	7,467.0	14,055.7	25,218.4	33,055.1
Start 3	7,439.0	13,693.0	26,066.6	33,987.1
Start 4	6,702.9	13,240.8	25,660.8	31,668.2
Start 5	8,670.8	14,147.9	26,173.4	33,271.5

Notes: per-percentage-point implied marginal cost scaled by the average market size (13,866 subscribers). λ -only component; the full marginal cost adds the firm–market–year cost shock $\hat{v}$.

The same conclusion emerges when translating the estimates into economically meaningful units. Table 18 reports the implied cost of increasing coverage by one percentage point in an average market. Although the underlying parameter estimates differ somewhat across solutions, the resulting costs are remarkably stable. The implied cost of a one-percentage-point expansion ranges from approximately \$6,700 to \$8,700 for UMTS rural markets and from \$25,000 to \$26,000 for CDMA markets, with the preferred specification lying comfortably within these intervals.

Table 19 examines the stability of the relative cost structure implied by the first-stage estimates. Because cost ratios are invariant to the monetary denomination of the cost function, they provide a scale-free summary of the economic content of the λ estimates that can be compared directly across starting values. The ratios are tightly clustered across the five multistart solutions: the CDMA–UMTS cost ratio ranges from 3.02 to 3.83 in rural areas and from 2.35 to 2.50 in urban areas, while the urban–rural differential ranges from 1.63 to 1.98 for UMTS and from 1.23 to 1.34 for CDMA. Two qualitative rankings hold in every solution. First, CDMA deployment is

Table 19

Relative Coverage Costs Across Standards and Terrain (Denomination-Free Ratios)

Lambda vector	CDMA/UMTS rural	CDMA/UMTS urban	Urban/rural UMTS	Urban/rural CDMA
Start 1 (= selected $\hat{\lambda}$)	3.083	2.504	1.650	1.340
Start 2	3.377	2.352	1.882	1.311
Start 3	3.504	2.482	1.841	1.304
Start 4	3.828	2.392	1.975	1.234
Start 5	3.019	2.352	1.632	1.271

Notes: Each cell reports a ratio of implied marginal costs of a one-percentage-point coverage expansion, constructed from the first-stage cost coefficients. Being ratios, the entries are invariant to any multiplicative rescaling of the cost function (e.g., the monetary denomination or the market-size normalization). Each row corresponds to one of the five multistart solutions of the MPEC estimation; Start 1 is the selected estimate used throughout the paper. The qualitative rankings, CDMA costlier than UMTS in both terrains, and urban costlier than rural for both standards, hold in every solution.

substantially more costly than UMTS deployment in both terrains, consistent with the narrower equipment vendor base of the CDMA ecosystem relative to the global GSM/UMTS supply chain. Second, urban expansion carries a cost premium over rural expansion for both standards, with the premium proportionally smaller for CDMA, whose rural deployment cost is already high. The fact that these relative magnitudes are nearly identical across solutions that differ visibly in their raw coefficients indicates that the data identify the relative cost structure robustly, independently of the level normalization at which any particular solution settles.

Overall, these results suggest that the model delivers stable and economically consistent implications even when initialized from substantially different starting values.

F.9 Numerical verification of identification.

Proposition 1 establishes the uniqueness of the equilibrium mapping from parameters to observables, while Proposition 2 establishes the identification through invertibility of the investment cost parameters. At the estimation stage, I verify numerically that the MPEC is well-conditioned at the converged solution: (i) the condition number of the constraint Jacobian is computed via singular value decomposition; (ii) the GMM moment Jacobian $G(\hat{\theta})$ is verified to have full column rank; and (iii) the information matrix $I(\hat{\theta}) = G(\hat{\theta})^\top W G(\hat{\theta})$ is confirmed to be invertible with smallest eigenvalue bounded away from zero. Together with the convergence to the same estimates across multiple starting values, these diagnostics provide empirical confirmation that the identification conditions in Proposition 1 and Proposition 2 are satisfied at the estimated parameter values.

Table 20

Numerical Verification of Identification

Diagnostic	Value
<i>(i) Constraint Jacobian SVD</i>	
Largest singular value	2.5243e+00
Smallest non-zero singular value	7.9229e-01
Condition number	3.1861e+00
Effective zero singular values	0 / 819
<i>(ii) Moment Jacobian G</i>	
Numerical rank	11 / 11
Full column rank	Yes
<i>(iii) Information matrix $I = G'WG$</i>	
Min eigenvalue	3.3959e-04
Max eigenvalue	5.4531e+01
Condition number	1.6058e+05
Invertible	Yes

Note: Local identification diagnostics for the MPEC.

F.10 Model Fit and Validation

To assess the validity of the structural estimates, this section compares the observed empirical moments from the data with the simulated moments generated by the model at the recovered parameter values. Overall, the model performs remarkably well in reproducing the targeted moments, capturing the key cross-sectional and time-series patterns for both consumer demand and firms' spatial network investment decisions.




Figure 12

Mean inside share: data vs. model

Figure 12 evaluates the fit of the demand model by comparing the mean inside market shares by provider and by standard. The simulated shares closely track the empirical data, replicating the aggregate market structure across the four national providers. Specifically, the model precisely captures the dominant market positions of Verizon Wireless and AT&T Mobility relative to Sprint and T-Mobile (Panel A). Furthermore, when aggregating these shares by technological standard, the simulated demand for both UMTS and CDMA networks matches the observed

data (Panel B). The absence of any systematic over- or under-prediction in these baseline shares confirms that the random coefficients logit demand system, combined with the recovered unobserved quality shocks, accurately reflects consumers' aggregate substitution patterns.




Figure 13
Mean total investment by standard: data vs. model

Turning to the supply side, Figure 13 demonstrates that the model captures the mean total coverage investment levels, correctly reflecting the higher average investment intensity of CDMA carriers compared to their UMTS counterparts. Moreover, Figure 14 compares observed and simulated investment moments across standards and years. The model closely matches average investment levels, including the higher investment intensity of CDMA2000 carriers, and reproduces the decline in coverage expansion as the 3G rollout matures.

Figures 15, 16, and 17 decompose this total investment into its spatial components, specifically, overlapping coverage expansion and roaming coverage expansion over time. The model successfully reproduces the qualitative time-series trajectories of both investment types across the 2015–2018 period, including the hump-shaped peak in overlapping investment in 2017 and the corresponding initial dip in new roaming coverage in 2016. While the model systematically, though slightly, underpredicts the absolute magnitude of overlapping investment across both standards (Figures 15 and 16) and proportionally overpredicts the extent of roaming coverage generated (Figures 15 and 17), the temporal dynamics and the relative differences across




Figure 14
Per-year moments: data vs model




Figure 15
Per-year overlapping and roaming Δ moments: data vs model

technological standards remain highly consistent with the empirical data. These deviations are modest relative to the overall scale of investment, confirming that the structural spatial matching function effectively captures the aggregate trade-offs firms face when allocating budgets between unique and shared geographical footprints.

G Additional Counterfactual Results and Robustness

This appendix reports additional counterfactual exercises that assess the robustness of the main welfare results to the treatment of prices and equilibrium selection. Table 21 holds national prices fixed at their observed levels, isolating the coverage channel. Table 22 reports the full counterfactual equilibrium with Bertrand–Nash prices, while Table 23 evaluates the sensitivity of the results to alternative equilibrium-selection rules.

Table 21 shows that the welfare gains persist when prices are held fixed. Total surplus net of investment increases by approximately \$4.51 million under all-UMTS and \$1.76 million under




Figure 16
Per-year total overlapping coverage moments by standard: data vs model




Figure 17
Per-year total roaming coverage moments by standard: data vs model

all-CDMA. Consumer surplus accounts for virtually all of the gains in this exercise, while producer surplus changes little. Investment responses are identical to those in the full counterfactual because prices are held fixed only after the investment equilibrium is determined.

Comparing Tables 21 and 22 shows that allowing firms to re-optimize prices mainly redistributes surplus between consumers and producers. Aggregate welfare changes only modestly across the two exercises, confirming that the welfare gains from standardization are driven primarily by the coverage response rather than by endogenous pricing.

Following Jia (2008), I bracket the counterfactual between the two extreme equilibria most favorable to the incumbent and to the fringe carriers, which bound the set of equilibrium outcomes; whereas Jia (2008) finds these extremes differ in a small share of markets, here they coincide exactly in every non-contested market-period, so the counterfactual is invariant to selection.

Specifically, Table 23 reports the counterfactual outcomes under alternative equilibrium-selection

Table 21
Detailed Counterfactual Outcomes at Fixed (Observed) Prices

	Levels			Difference vs. Multiple	
	Multiple	All-UMTS	All-CDMA	UMTS	CDMA
<i>Social welfare</i>					
Consumer surplus	50438782	55265042	52593718	+4826260	+2154936
Producer surplus (variable)	5314560	5012655	5325423	-301904	+10864
Total surplus	55753342	60277698	57919142	+4524356	+2165800
Investment expenditure	1160748	1174513	1562505	+13765	+401757
Total surplus net of investment	54592594	59103184	56356637	+4510591	+1764043
Δ per person-year					
Δ CS / person-year	–	+1.313	+0.586		
Δ TS / person-year	–	+1.231	+0.589		
<i>Producer surplus by carrier</i>					
AT&T Mobility	2028721	1252952	1684740	-775769	-343981
Sprint	608092	1359732	856200	+751640	+248109
T-Mobile	490183	662168	445720	+171985	-44463
Verizon Wireless	2187564	1737803	2338763	-449761	+151199
<i>Investment</i>					
y_{tot} total	23.1502	36.7108	23.7018	+13.5606	+0.5517
y_E total	16.2059	10.2498	8.1496	-5.9561	-8.0563
y_C total	6.9443	26.4610	15.5522	+19.5167	+8.6079
AT&T Mobility: mean y_{tot}	0.0266	0.0332	0.0340	+0.0065	+0.0074
Sprint: mean y_{tot}	0.0451	0.0772	0.0414	+0.0321	-0.0037
T-Mobile: mean y_{tot}	0.0367	0.0802	0.0309	+0.0435	-0.0057
Verizon Wireless: mean y_{tot}	0.0610	0.0782	0.0671	+0.0172	+0.0061
Prices fixed at observed values in all columns by construction.					

Note: National prices held at observed values (coverage channel only). Units and certification as in Table 5 in the paper. Results shown for the status-quo selection; all selection rules coincide (unique equilibria).

rules that favor the status quo, incumbent carriers, or fringe carriers when multiple equilibria are detected. The aggregate welfare and investment results are identical across all three rules. The counterfactual algorithm identifies a unique verified pure-strategy equilibrium in all non-contested market-periods; the remaining contested cells, 31 under all-UMTS and 34 under all-CDMA, are resolved identically under all three selection rules. The resulting welfare and investment comparisons are therefore invariant to equilibrium selection.

Table 22
Detailed Counterfactual Outcomes at Re-Equilibrated Prices

	Levels			Difference vs. Multiple	
	Multiple	All-UMTS	All-CDMA	UMTS	CDMA
<i>Social welfare</i>					
Consumer surplus	50438782	53718714	51050770	+3279932	+611988
Producer surplus (variable)	5314560	6488542	6671273	+1173982	+1356713
Total surplus	55753342	60207255	57722043	+4453913	+1968701
Investment expenditure	1160748	1174513	1562505	+13765	+401757
Total surplus net of investment	54592594	59032742	56159538	+4440148	+1566945
<i>Δ per person-year</i>					
Δ CS / person-year	–	+0.893	+0.167		
Δ TS / person-year	–	+1.212	+0.536		
<i>Producer surplus by carrier</i>					
AT&T Mobility	2028721	1569646	2212182	-459075	+183461
Sprint	608092	1844038	1096266	+1235946	+488174
T-Mobile	490183	960485	619174	+470302	+128991
Verizon Wireless	2187564	2114373	2743651	-73192	+556086
<i>Investment</i>					
y_{tot} total	23.1502	36.7108	23.7018	+13.5606	+0.5517
y_E total	16.2059	10.2498	8.1496	-5.9561	-8.0563
y_C total	6.9443	26.4610	15.5522	+19.5167	+8.6079
AT&T Mobility: mean y_{tot}	0.0266	0.0332	0.0340	+0.0065	+0.0074
Sprint: mean y_{tot}	0.0451	0.0772	0.0414	+0.0321	-0.0037
T-Mobile: mean y_{tot}	0.0367	0.0802	0.0309	+0.0435	-0.0057
Verizon Wireless: mean y_{tot}	0.0610	0.0782	0.0671	+0.0172	+0.0061
<i>Mean national price by carrier</i>					
AT&T Mobility	3.622	3.542	3.673	-0.079	+0.051
Sprint	3.202	4.174	3.717	+0.973	+0.516
T-Mobile	3.351	4.366	4.092	+1.015	+0.742
Verizon Wireless	3.763	3.984	4.375	+0.222	+0.612

Note: Bertrand–Nash prices at the certified investment allocation. Units and certification as in Table 5 in the paper. Results shown for the status-quo selection; all selection rules coincide (unique equilibria).

Table 23
Counterfactual outcomes under alternative equilibrium-selection rules

	All-UMTS			All-CDMA		
	Status quo	Favor dominant	Favor fringe	Status quo	Favor dominant	Favor fringe
Δ Consumer Surplus	+3279932	+3279932	+3279932	+611988	+611988	+611988
Δ Producer Surplus	+1173982	+1173982	+1173982	+1356713	+1356713	+1356713
Δ Social Welfare	+4453913	+4453913	+4453913	+1968701	+1968701	+1968701
Δy_{tot} mean	+0.0248	+0.0248	+0.0248	+0.0010	+0.0010	+0.0010
Δy_E mean	-0.0109	-0.0109	-0.0109	-0.0148	-0.0148	-0.0148
Δy_C mean	+0.0357	+0.0357	+0.0357	+0.0158	+0.0158	+0.0158
Contested cells	31	31	31	34	34	34

Note: Δ shown versus the Multiple-standard baseline, at the Bertrand–Nash price equilibrium given the certified investment allocation; welfare in model profit units, investment in coverage units (means per market-year, one unit = 100 pp). The three selection rules coincide because the counterfactual equilibrium is unique in every non-contested market-period (no cell admits multiple verified pure-strategy equilibria); audited maximum absolute difference across rules, over both aggregates and the per-cell allocation: 0.00e+00. Contested cells (31 UMTS, 34 CDMA) are resolved identically under all rules.

H Appendix: Base stations coverage measure

Mobile telephony is called “cellular” in the United States due to the practice of dividing space up into discrete “cells” served by separate base stations. Each grouping or “cluster” of base stations has access to all the firm’s licensed frequency. If consumers move out of the range of a cluster’s cell into a new area, they are simply transferred to the cell that covers that area and its assigned frequency. In this way, a firm can reuse a limited amount of frequency, and this innovation made mass adoption of mobile phones possible.

How far a signal can go from a tower depends on several factors, such as the geographical configuration of the market and the technology standard deployed. Since I do not observe the specific geographical characteristics or the number of towers built in a given market, to measure the geographical area covered by a new base station I assume completely flat terrain. The most efficient base-station deployment then has base stations at the centers of identical regular hexagons that tile the space completely. I further assume that a new base station is built at the center of each hexagon and that each station increases the covered area equally, i.e., there is no decreasing marginal coverage from additional stations.

Since hexagons are regular geometric figures, to determine the area covered by a base station I divide each hexagon into 12 congruent right triangles with base a , height $b = \frac{\sqrt{3}}{2}c$, and hypotenuse c , where c equals the signal radius. According to engineers and industry experts, the maximum distance a signal can be transmitted from a base station ranges between 35.5 and 72.5 km, depending on terrain and the technology deployed. Taking an average distance of 54 km gives $a = 27$, $b = 46.77$, and $c = 54$ km, so the area of each right triangle is

$$Area_{abc} = \frac{a \cdot b}{2} = \frac{27 \cdot 46.77}{2} = 631.33 \text{ km}^2. \quad (24)$$

The hexagon area is then $Area_{abc} \times 12 = 7,576 \text{ km}^2$. Transforming this into percentage points using the area of each DMA and averaging across markets, I obtain that a new base station leads to a constant increase in coverage of about 17.5% across markets on average.

To translate the engineering estimates into a cost per percentage point of coverage, I combine this coverage increment with the cost of a macro site. The implied coverage gain from a single site depends on the assumed signal radius. At the upper end of the 35.5–72.5 km range a site covers a large share of a market and the implied cost per percentage point is low, whereas at the lower end each site covers far less area and the cost per point is correspondingly higher. Taking the generous midpoint radius of 54 km (the 17.5% increment above) and combining it with the \$266,000–\$285,000 cost of a macro site over a 15-year horizon ([Johansson et al.](#),

2004; Claussen et al., 2008; Björkegren, 2019) yields a one-time cost of roughly \$16,000 per percentage point. This figure should be read as a *lower bound*: it credits each site with the coverage it would deliver only under the most favorable propagation assumptions and under no decreasing marginal coverage. Using instead the lower end of the cited range (a 35.5 km radius), a site covers closer to 7.6% of a market, raising the implied engineering cost to roughly \$36,000 per percentage point. The engineering benchmark therefore spans approximately \$16,000 to \$36,000 per percentage point of coverage, with the true marginal cost likely toward the upper end once the decreasing marginal coverage of additional sites in already-served markets is taken into account.

Bibliography

- BERRY, S. T. (1994): “Estimating discrete-choice models of product differentiation,” *The RAND Journal of Economics*, 242–262.
- BJÖRKEGREN, D. (2019): “The adoption of network goods: Evidence from the spread of mobile phones in Rwanda,” *The Review of Economic Studies*, 86, 1033–1060.
- CLAUSSEN, H., L. T. HO, AND L. G. SAMUEL (2008): “An overview of the femtocell concept,” *Bell Labs Technical Journal*, 13, 221–245.
- DUBÉ, J.-P., J. T. FOX, AND C.-L. SU (2012): “Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation,” *Econometrica*, 80, 2231–2267.
- GOOLSBEE, A. AND A. PETRIN (2004): “The consumer gains from direct broadcast satellites and the competition with cable TV,” *Econometrica*, 72, 351–381.
- JIA, P. (2008): “What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry,” *Econometrica*, 76, 1263–1316.
- JOHANSSON, K., A. FURUSKAR, P. KARLSSON, AND J. ZANDER (2004): “Relation between base station characteristics and cost structure in cellular systems,” in *2004 IEEE 15th International Symposium on Personal, Indoor and Mobile Radio Communications (IEEE Cat. No. 04TH8754)*, IEEE, vol. 4, 2627–2631.
- SU, C.-L. AND K. L. JUDD (2012): “Constrained optimization approaches to estimation of structural models,” *Econometrica*, 80, 2213–2230.