Price Cap vs. Ad Valorem Subsidies: Selection, Pricing, and Cross-Subsidization in the FCC's Rural Health Care Program

Ram Sewak Dubey ⓡ Maysam Rabbani ⓡ Rodrigo Pinto*

May 29, 2026

Abstract

We evaluate the FCC's Rural Health Care subsidy mechanisms using administrative data and an extension of two-way fixed effects. The original price-cap mechanism lacked cost-containment incentives. A 2014 reform introduced an ad valorem mechanism with 35% cost-sharing. We find that switching from price-cap to ad valorem reduces subsidy payment by 59%. However, the reform introduced consortium applications. We document how this inadvertently enabled cross-subsidization from eligible to ineligible members, inflating outlays by 13% ($70M annually). We recommend discontinuing both the price cap and the consortium option, leaving individual ad valorem as the sole mechanism.

Keywords: `Policy Evaluation`, `Subsidies`, `Telecommunication`.
JEL codes: D04, D22, H25, I38, L96

*Economics Department, Feliciano School of Business, Montclair State University, Montclair, NJ, USA (Dubey, Rabbani); Department of Economics, University of South Florida, FL, USA (Pinto). We are grateful for the interaction with Thomas E Cone, the seminar participants at Montclair State University, University of South Florida, New York State Economics Association, and Southern Economics Association. Author names have been randomized. The official randomization link may be found here.

1 Introduction

The Rural Health Care Program (RHC) spends over $700 million per year subsidizing internet connectivity for rural health care providers (HCPs) across the United States (US). Under the program's original price-cap scheme, a rural hospital pays the urban benchmark price regardless of what the Internet Service Provider (ISP) charges, and the government covers the difference. The hospital's out-of-pocket cost is therefore fixed by design, leaving no incentive to contain costs. In 2014, policymakers introduced an ad valorem alternative in which hospitals pay 35% of the negotiated price, restoring their incentive to bargain. The reform also permitted *consortium* applications that mix eligible and ineligible HCPs, inadvertently letting the former cross-subsidize the latter and extract additional subsidies.

This paper provides the first comprehensive economic evaluation of the RHC. We exploit the 2014 introduction of the Healthcare Connect Fund (HCF), which offered every price-cap participant the option of staying on $\mathcal{P}_1$ or switching to the ad valorem mechanism ($\mathcal{P}_2$) or its consortium variant ($\mathcal{P}_2^c$). The 2013–2014 year pair is a clean natural experiment. $\mathcal{P}_2$ and $\mathcal{P}_2^c$ became available simultaneously in 2014, leaving only the $\mathcal{P}_1 \to \mathcal{P}_2$ and $\mathcal{P}_1 \to \mathcal{P}_2^c$ switching margins. We extend the analysis to seven additional year pairs through 2020–2021 and combine the eight year-pair estimates by inverse-variance weighting to obtain headline magnitudes on a broader sample. Switching from $\mathcal{P}_1$ to $\mathcal{P}_2$ reduced internet price by 58%, subsidy spending by 59%, and HCP net cost by 31%, consistent with the prediction that restoring demand elasticity lowers the monopolist's equilibrium markup. Switching from $\mathcal{P}_1$ to $\mathcal{P}_2^c$ reduced subsidy spending by only 25% and *increased* HCP net cost by 74%. That HCPs voluntarily accepted a materially costlier mechanism is hard to explain without latent incentives such as cross-subsidization from eligible to ineligible consortium members.

These findings speak to a general design vulnerability. Whenever a subsidy program permits joint applications by eligible and ineligible participants, the group can strategically reallocate costs to maximize subsidy extraction. By inflating the reported price for eligible members and deflating it for ineligible ones, a consortium extracts higher subsidies without changing the ISP's total revenue. The gains are bounded only by enforcement, which in this program appears weak. GCI Communication Corp paid a $42.1 million settlement for price misreporting (Appendix A.3), a penalty equal to just 3.4% of the subsidies the firm collected over the observed years.

Our theoretical model derives testable implications from monopoly pricing under each mechanism. Under the price cap, the HCP pays the benchmark price regardless of what is billed, so demand is perfectly inelastic and the monopolist inflates the billed price until the marginal enforcement penalty equals the marginal revenue from overcharging. Under the ad valorem subsidy, the HCP pays a fixed percentage of the billed price, restoring price-elastic demand and a standard monopoly markup. We show that under empirically relevant conditions the ad valorem mechanism dominates the price cap in consumer price, quantity, HCP expenditure, and government outlays. For consortia, we model cross-subsidization as a multiplicative distortion applied to the eligible member's price, with detection risk convex in the distortion. The optimal distortion is inverted U-shaped in the ratio of ineligible to eligible revenue, since small ratios offer little revenue to shift and large ratios trigger steep enforcement penalties.

We use the full population of RHC subsidy requests recorded in Federal Communications Commission (FCC)-mandated administrative data. We extend the conventional Two-Way Fixed Effects (TWFE) specification to accommodate two simultaneous program-share measures rather than a single binary indicator. For each HCP we measure these shares as the fraction of bandwidth allocated to $\mathcal{P}_2$ and the fraction allocated to $\mathcal{P}_2^c$. We test the parallel-trends assumption directly using 2012 data. Because HCF was not available before 2013, the 2012→2013 panel serves as a clean pre-period placebo. We also assess sensitivity of the difference-in-difference (DiD) estimates to a parameterized departure from parallel trends following Aryal et al. (2025) and Manski and Pepper (2018).

We supplement these estimates with double/debiased machine learning (DML, Chernozhukov et al., 2018) to allow for nonlinear covariate relationships and report robustness checks across HCP types, service types, and sample restrictions. Switching remains a voluntary decision, so our estimates, which target the average treatment effect on the treated, may reflect selection into mechanisms. We address this concern through several complementary approaches: an institutional argument that the three programs differ only in their reimbursement formula, a logit analysis showing that a single mechanical cost signal dominates the switching decision, a sensitivity analysis that allows proportional departures from parallel trends (Aryal et al., 2025), and a theory-derived test of cross-subsidization that does not rely on parallel trends.

The evidence for cross-subsidization rests on four complementary sources. First, the theoretical model predicts that consortia inflate eligible members' prices whenever detection is imperfect. Second, the empirical estimates show that switching to $\mathcal{P}_2^c$ increased HCP net cost, a pattern that is difficult to rationalize since the ad valorem rate is identical under $\mathcal{P}_2$ and $\mathcal{P}_2^c$. Third, a local linear regression of cost distortion on consortium composition reveals an inverted U-shape, exactly as the model predicts. Fourth, FCC enforcement records, including settlements for price misreporting in consortium applications, provide direct institutional corroboration of the mechanism. No single piece of evidence is conclusive on its own. Together, they form a coherent case suggesting that cross-subsidization is economically significant and consistent with the theoretical predictions.

The cross-subsidization channel we identify requires three conditions: subsidies tied to a reported metric that participants can manipulate, a group structure that permits internal reallocation across eligible and ineligible members, and enforcement too weak to deter distortion at the margin. These conditions are not unique to the studied programs. The Universal Service Fund (USF)'s own E-Rate program satisfies all three. E-Rate allows consortia to comprise eligible and ineligible members, and the subsidy ranges from 20% to 90% based on school demographics. E-Rate's effectiveness has been challenged on multiple fronts (Goolsbee and Guryan, 2006; Hazlett, Schwall, and Wallsten, 2019). The cross-subsidization channel our framework predicts has not been empirically tested in E-Rate, making it a direct out-of-sample implication. More broadly, similar conditions may arise in health care, housing, and energy programs that permit joint applications by heterogeneous participants.

Our paper contributes to three literatures. First, we advance the empirical literature on subsidy design and program effectiveness (Caliendo and Künn, 2011; Eriksson, Kaserman, and Mayo, 1998; Fairlie and Robinson, 2013; Goolsbee and Guryan, 2006; Gruber and Washington, 2005; Mendez, Molnar, and Savage, 2021) by providing causal evidence on how price-cap versus ad valorem subsidies affect negotiated prices and participant costs. Second, we identify a novel form of regulatory leakage, namely cross-subsidization from eligible to ineligible consortium members, complementing the literature on unintended consequences of government interventions (Cicala, Lieber, and Marone, 2019; Fremeth, Richter, and Schaufele, 2018; Vernon, 1979). Our channel is the mirror image of Rotemberg (2019). Whereas subsidies to eligible firms crowd out ineligible rivals in his setting, ineligible entities benefit *from within* the subsidized group in ours. The mechanism also parallels the upcoding evidence in Medicare Advantage (Geruso and Layton, 2020) and the bid-formula gaming documented in Medicare Part D (Decarolis, Polyakova, and Ryan, 2020), with the distinguishing feature that gaming here is collective across consortium members rather than individual. Third, on the methodological side, we extend TWFE to handle two simultaneous treatment-share variables and combine it with DML (Aryal et al., 2025; Chernozhukov et al., 2018; Manski and Pepper, 2018).

The paper proceeds as follows. Section 2 describes the institutional environment. Section 3 presents the theoretical model. Section 4 develops the econometric framework. Section 5 reports empirical results. Section 6 discusses policy implications. Section 7 concludes.

2 Institutional environment[1]

The FCC was established in 1934 with a mandate to make rapid, efficient, nationwide communications available to all people in the US at reasonable charges (47 U.S.C. § 151). The Telecommunications Act of 1996 created the USF, an umbrella of four sister programs that target distinct beneficiary groups (Figure 1): the High Cost Program reimburses telecommunication firms serving high-cost areas, E-Rate subsidizes internet access for schools and libraries, Lifeline supports low-income consumers, and the RHC subsidizes internet access for eligible HCPs.

The FCC appointed the Universal Service Administrative Company (USAC), a private nonprofit, to administer the USF. Instead of drawing on the federal budget, USAC levies a quarterly "contribution factor" on telecommunication companies, which firms are permitted to itemize and pass through to end users. The contribution factor has risen persistently from 3.14% in 1998 to 37.6% in 2026 (Figure A1). The four USF programs together cost \$8.6 billion in 2024 (USAC, 2024), borne by internet users. Independent audits and policymakers have expressed concern that pro-

[1] Appendix A provides a more detailed version of this section.

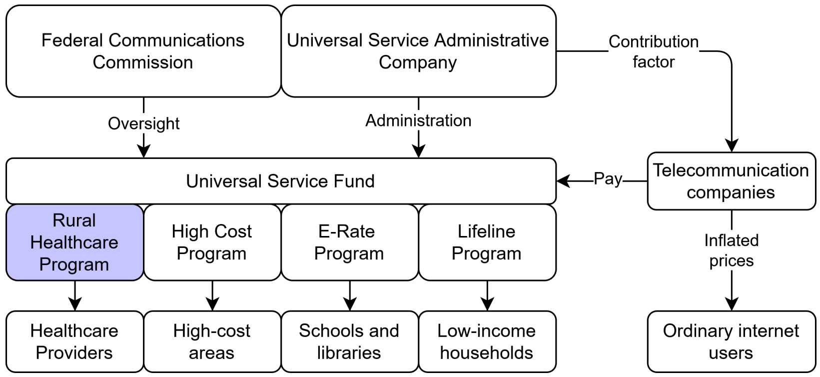

Figure 1: The regulatory structure and key players.

Notes: The Rural Health Care Program is one of four Universal Service Fund subsidy programs administered by the Universal Service Administrative Company under FCC oversight. The fund is financed by a tax on telecommunication carriers that is passed through to end users.

gram costs are inflated by waste, fraud, and abuse, a concern that our cross-subsidization findings corroborate. The financing structure may also generate a positive feedback loop. Rising program costs raise the contribution factor, raising internet prices, expanding reliance on subsidies, and further inflating costs (Hazlett, Schwall, and Wallsten, 2019). A back-of-the-envelope calculation in Appendix A suggests the contribution factor has raised average internet prices in the US by 7.2%, pushing 1.3–3.4% of users into subsidy dependence.

The other three USF programs have been extensively studied: Lifeline (Ackerberg et al., 2014; Conkling, 2018; Lyons, 2023; Mendez, Molnar, and Savage, 2021; Wallsten, 2016; Ward and Woroch, 2010), E-Rate (Goolsbee and Guryan, 2006; Greenstein, 2020; Hazlett, Schwall, and Wallsten, 2019), and High Cost (Berg, Jiang, and Lin, 2011a,b; Boik, 2017; Hanna, Savage, and Wimmer, 2025; Wallsten, 2011). The RHC has received less attention. Rabbani (2024b) documents that HCP internet became faster and cheaper between 2014 and 2020, and Rabbani (2024a) estimates that $\mathcal{P}_1$ plans cost 132–179% more than $\mathcal{P}_2$ plans on similar service, with an implied \$143 million in annual taxpayer savings from discontinuing $\mathcal{P}_1$. Our paper provides the first causal evaluation grounded in a theoretical model and quantifies the consortium channel that prior work did not address.

Our analysis focuses on the RHC program, whose institutional timeline is shown in Figure 2. Established in 1997, the program initially used a single price-cap subsidy mechanism known as the Telecommunications Program ($\mathcal{P}_1$). At rural internet price p, if an HCP could demonstrate that the same plan cost p_u in a nearby urban area, the HCP would pay p_u while the program covered $p - p_u$. This design is flawed economically because, at a fixed p_u, $\mathcal{P}_1$ insulates HCPs from marginal changes in p. Demand becomes perfectly inelastic, eliminating the incentives for cost minimization (Rabbani, 2024a).

The Healthcare Connect Fund (HCF) was piloted in 2013 and fully implemented in 2014. Rather than benchmarking against urban prices, HCF uses an ad valorem subsidy that covers 65% of the negotiated price, leaving the HCP a 35% copayment. This proportional cost-sharing

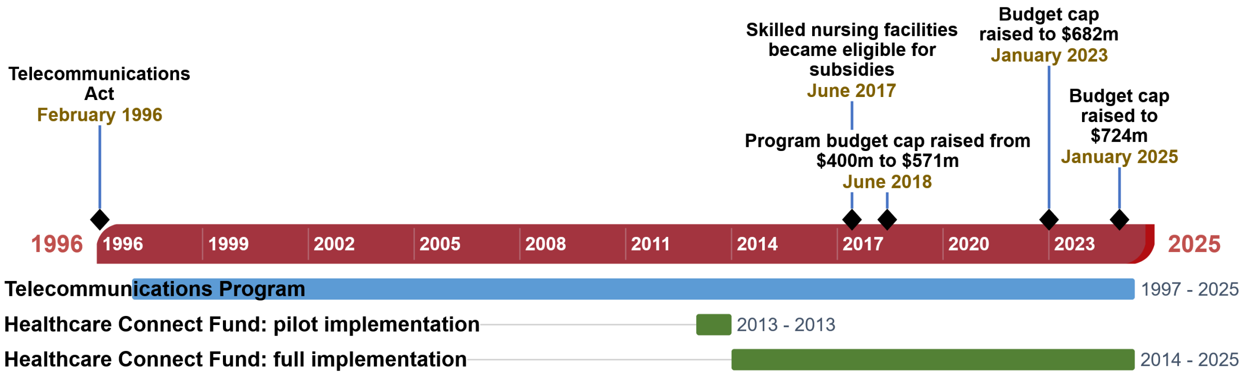

Figure 2: A timeline of program implementation.

Notes: Timeline of the Rural Health Care Program and its two subsidy mechanisms, the Telecommunications Program ($\mathcal{P}_1$) and the Healthcare Connect Fund ($\mathcal{P}_2$ and $\mathcal{P}_2^c$).

introduces efficiency incentives. HCPs gain $0.35 for every $1.00 saved and lose $0.35 for every $1.00 wasted. Under HCF, participants may seek subsidies individually ($\mathcal{P}_2$) or jointly through a consortium that requests subsidies on behalf of its members ($\mathcal{P}_2^c$). The consortium option extends benefits to urban HCPs if the consortium is majority-rural, i.e., if more than half of its members are in rural areas. For example, a large urban hospital that forms a consortium with two small rural clinics is majority-rural by headcount, and all three members would receive subsidies.

An HCP is eligible for subsidies if (1) it is a non-profit or public entity, (2) it is rural or part of a majority-rural consortium, and (3) it belongs to a qualifying entity type (listed in Appendix A.2). Eligible HCPs may choose $\mathcal{P}_1$, $\mathcal{P}_2$, or $\mathcal{P}_2^c$, with the exception that eligible urban HCPs can only access subsidies via $\mathcal{P}_2^c$. HCPs may switch programs upon subsidy renewal. HCPs that fail any of the three criteria are ineligible. Ineligible HCPs receive no subsidies and do not count toward majority-rural status, but they may join a consortium to seek lower internet prices through collective bargaining.

When a consortium submits a subsidy request, it enters a competitive bidding process in which ISPs place bids for providing internet services to all eligible and ineligible members. The auction winner is not the lowest bid. Instead, the HCP chooses the bid it deems most desirable based on discretionary criteria (see Table A2). ISPs therefore compete to make their bid most appealing to the HCP, not to offer the lowest price. We hypothesize that this enables cross-subsidization: the winning bid inflates prices for eligible members and deflates prices for ineligible members, raising subsidy outlays. For every $1 cross-subsidized, the consortium collects an additional $0.65 in subsidies. This can be revenue-neutral to the ISP, or the ISP may capture some of the proceeds. Decarolis, Polyakova, and Ryan (2020) document an analogous bid-formula response in Medicare Part D, and Appendix A.3 documents the RHC's institutional conditions and supportive anecdotal evidence.

Figure 3 traces program activity from 2013 to 2021. Panel D shows the migration of HCPs: $\mathcal{P}_1$ shrinks from 1,477 to 164 HCPs, while $\mathcal{P}_2^c$ grows from 238 in 2014 to 4,061 by 2021 and overtakes

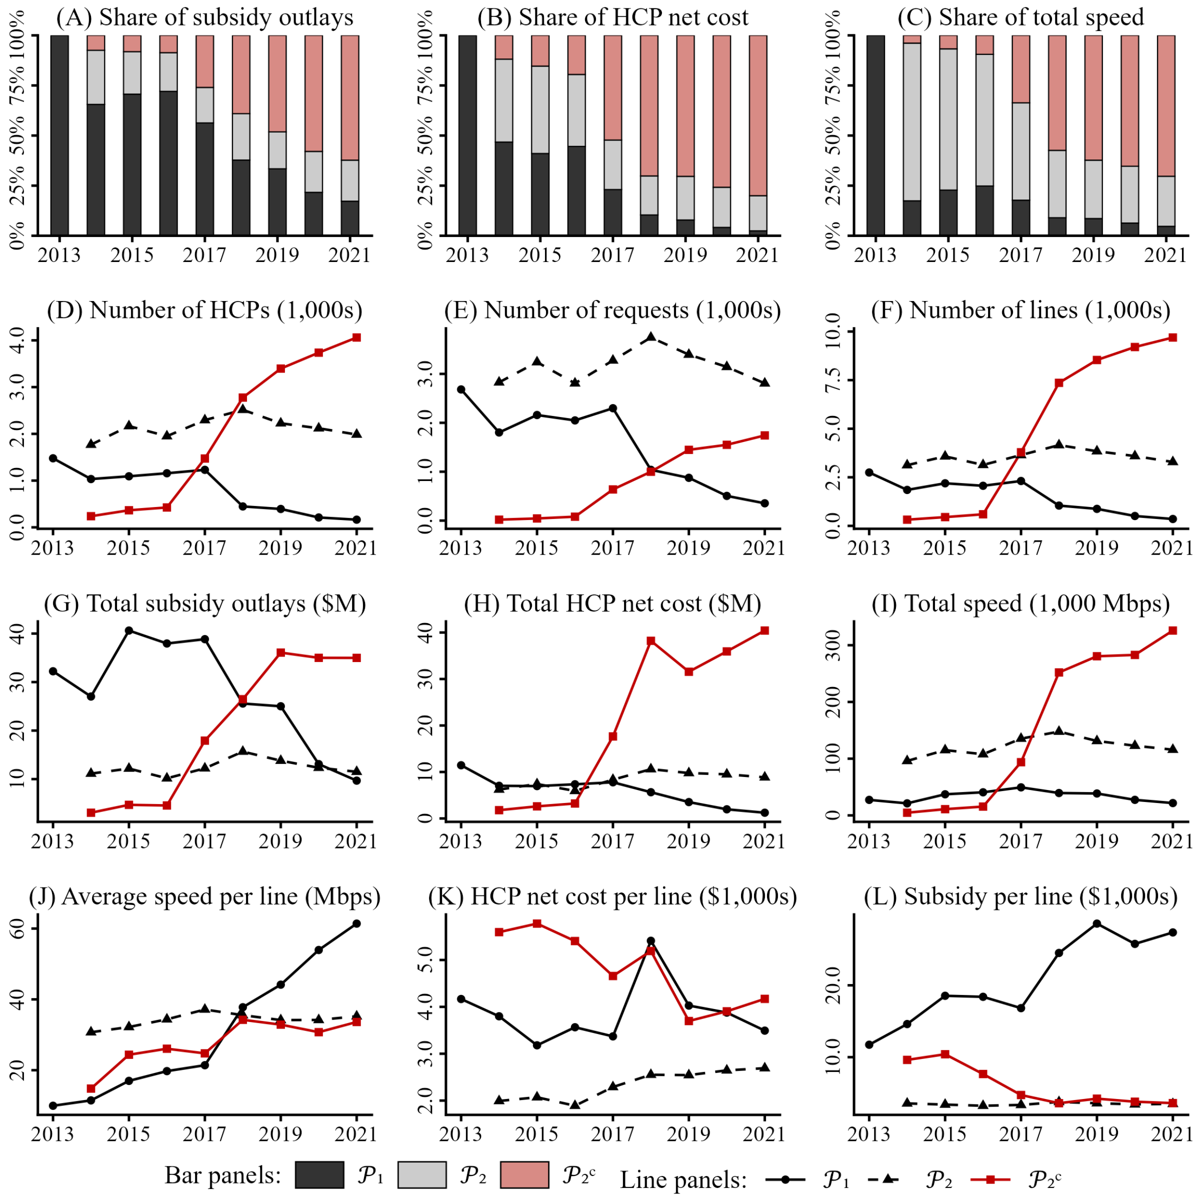

Figure 3: Year-by-year program evolution, 2013–2021.

Notes: Each panel plots an annual aggregate by program on the baseline analysis sample. Row 1 reports shares (summing to 100% within year); rows 2–4 report levels. $\mathcal{P}_2$ and $\mathcal{P}_2^c$ are first observed in 2014.

$\mathcal{P}_2$ in 2018. Panels G–H show the corresponding shift in dollar flows: by 2021, $\mathcal{P}_2^c$ absorbs three times more in subsidy outlays than $\mathcal{P}_1$. Panel K shows that HCP net cost per line is systematically lowest under $\mathcal{P}_2$ and highest under $\mathcal{P}_2^c$ throughout, an early indication that the two ad valorem variants deliver very different value despite the same nominal cost-sharing rate. Widespread $\mathcal{P}_2^c$ participation began after 2017, when regulatory changes expanded eligibility.[2]

[2]The first event was in June 2017, when skilled nursing facilities became eligible for subsidies, opening the program to many new HCPs. The second event was a budget adjustment. The program's annual budget was capped at $400 million at its inception in 1997, but the cap was first exceeded in 2016 after holding for nearly two decades, straining program finances. In June 2018, the FCC began indexing the cap to inflation for the first time in the program's

The first switching opportunity arrived in 2014. HCPs who chose $\mathcal{P}_2$ that year opted for higher speeds, yet the minority that stayed on $\mathcal{P}_1$ generated more than half of the RHC's cost burden despite purchasing a small share of total speed. A simple cost-minimization argument explains the pattern. In 2013, all HCPs were on $\mathcal{P}_1$. $\mathcal{P}_2$ requires the HCP to pay a 35% copayment, so an HCP prefers $\mathcal{P}_2$ whenever its HCP net cost ratio under $\mathcal{P}_1$ exceeds 35%. A low ratio implies that the rural rate is far above the urban benchmark, suggesting an overpriced plan (Rabbani, 2024a). The 2014 reform therefore induced a selective migration in which cost-efficient HCPs moved to $\mathcal{P}_2$ while the most overpriced plans remained in $\mathcal{P}_1$. The same pattern persists in later years: $\mathcal{P}_1$'s subsidy per line climbs from $14,606 in 2014 to $27,355 in 2021 (Panel L).

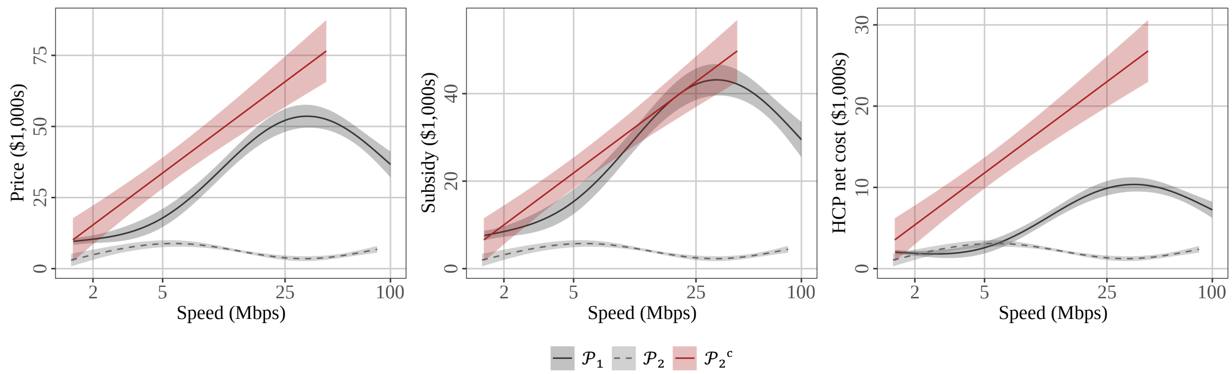

Figure 4: Price measures by program, 2014.

Notes: Generalized Additive Model fits at the request level using 2014 data, with 95% confidence ribbons. Each panel reports a dollar outcome (price, subsidy, or HCP net cost) against ln(speed). $\mathcal{P}_1$ dark gray, $\mathcal{P}_2$ light gray, $\mathcal{P}_2^c$ red.

Figure 4 compares price, subsidy, and HCP net cost across speed levels using a Generalized Additive Model (GAM), a semi-parametric method that accommodates nonlinear price–speed relationships. Each fitted curve is shown with a 95% confidence interval. Plans subsidized through $\mathcal{P}_2$ are cheaper than those in $\mathcal{P}_1$, consistent with the cost-effectiveness hypothesis. Plans in $\mathcal{P}_2^c$ are generally more expensive than those in $\mathcal{P}_2$, consistent with the cross-subsidization hypothesis. We formalize both hypotheses in Section 3.

An HCP may hold multiple active lines of subsidy simultaneously, especially if it operates across multiple buildings or locations. Many HCPs hold tens or hundreds of lines at once, often spread across $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_2^c$. Figure A3 shows how the resulting program-share composition evolves from 2013 to 2021. HCPs start at the $\mathcal{P}_1$ vertex in 2013. In 2014, migration is mostly to $\mathcal{P}_2$ or to $\mathcal{P}_1$–$\mathcal{P}_2$ mixes, with little movement to $\mathcal{P}_2^c$. Substantial $\mathcal{P}_2^c$ adoption emerges only after 2018, and by 2021 HCPs concentrate near $\mathcal{P}_2$, $\mathcal{P}_2^c$, or mixes of the two.

history, and the budget has since risen to $724 million in 2025. FCC: Rural Health Care Program, FCC: Report and Order, WC Docket No. 17-310

3 A monopoly model of subsidy design

This section develops a monopoly model linking subsidy design to equilibrium prices, quantities, and government outlays. The HCP is the consumer of a monopolistic ISP. Under a price-cap regime ($\mathcal{P}_1$), consumer demand becomes locally inelastic once the cap binds, severing the link between billed prices and quantities and allowing the monopolist to inflate bills against the government budget. Under an ad valorem subsidy ($\mathcal{P}_2$), the consumer retains proportional exposure to the posted price and stays sensitive to price changes, which constrains the monopolist's markup. When a consortium of eligible and ineligible members operates under $\mathcal{P}_2^c$, it can inflate prices for eligible members and use the additional subsidies to reduce costs for ineligible ones. This unintended consequence erodes the price discipline that the ad valorem mechanism was designed to restore.

3.1 Environment

A monopolist ISP supplies a homogeneous good at constant marginal cost $c > 0$. Demand $D : \mathbb{R}_+ \to \mathbb{R}_+$ satisfies $D'(p) < 0$ wherever $D(p) > 0$ and determines the quantity $Q = D(p_c)$ at HCP net cost p_c. Revenue $p\ D(p)$ is strictly concave. Without subsidies, the monopoly price p^{no} is the unique solution to the first-order condition:

$$D(p^{no}) + (p^{no} - c)\, D'(p^{no}) = 0. \tag{1}$$

Equivalently, the Lerner index gives:

$$\frac{p^{no} - c}{p^{no}} = \frac{1}{\varepsilon_D(p^{no})},$$

where $\varepsilon_D(p) \equiv -\frac{pD'(p)}{D(p)}$ is the price elasticity of demand. Under perfect competition, $p_c = c$ and $Q = D(c)$. [3]

3.2 Price-cap subsidy ($\mathcal{P}_1$)

Under $\mathcal{P}_1$, the ISP posts a billed price $p \geq 0$. The consumer pays at most a regulated cap $\bar{p} > 0$, so the HCP net cost is $p_c^{cap}(p) = \min\{p, \bar{p}\}$. When $p > \bar{p}$, the government reimburses the gap per unit:

$$G^{cap}(p) = (p - \bar{p})_+\, D\big(p_c^{cap}(p)\big), \tag{2}$$

where $(x)_+ \equiv \max\{x, 0\}$. Excess billing is subject to an expected penalty $\alpha\, \Phi(\delta)$ on the price deviation $\delta \equiv (p - \bar{p})_+$, where $\alpha \in (0, 1]$ is the probability of being audited and Φ is a penalty function satisfying $\Phi(0) = 0$, $\Phi'(0) = 0$, $\Phi'(\delta) > 0$ for $\delta > 0$, and $\Phi'' > 0$. The condition $\Phi'(0) = 0$ ensures that small deviations are not deterred, while the strict convexity of Φ ensures that marginal penalties escalate with the degree of overcharging. For tractability, we adopt the quadratic specification $\Phi(\delta) = \frac{\gamma}{2}\delta^2$ ($\gamma > 0$).

The economic mechanism is a *discontinuity in the residual demand elasticity.* When the cap does not bind ($p \leq \bar{p}$), the monopolist faces the standard downward-sloping demand curve and

[3] Although we present results for a general demand function, Appendix B.4 provides closed-form expressions under the linear specification $D(p_c) = a - bp_c$ ($a, b > 0$, $a > bc$).

markups are disciplined by the demand response. Once p crosses $\bar{p}$, the consumer price is pinned at $\bar{p}$, demand freezes at $D(\bar{p})$, and the firm's residual demand becomes perfectly inelastic in the billed price. The only force restraining further price inflation is the enforcement technology (α, Φ).

The cap binds whenever $\bar{p} < p^{no}$.[4] When it does, the HCP pays $p_c^{cap} = \bar{p}$ and demand is pinned at $Q^{cap} = D(\bar{p})$. The HCP expenditure is $E^{cap} = \bar{p}\, D(\bar{p})$, and the firm's problem reduces to:

$$\max_{p \geq \bar{p}} \, (p - c)\, D(\bar{p}) - \alpha \Phi(p - \bar{p}). \tag{3}$$

Proposition P.1 (Billed-price under $\mathcal{P}_1$)**.** Suppose the cap binds. Then the equilibrium billed price is

$$p^{cap} = \bar{p} + (\Phi')^{-1}\biggl(\frac{D(\bar{p})}{\alpha}\biggr), \tag{4}$$

which is increasing in $D(\bar{p})$ and decreasing in α. Under the quadratic penalty,

$$p^{cap} = \bar{p} + \frac{D(\bar{p})}{\alpha\gamma}, \qquad G^{cap} = \frac{D(\bar{p})^2}{\alpha\gamma}. \tag{5}$$

Proof. See Appendix B.1. □

The result exposes the fundamental weakness of price-cap subsidies. The marginal government payment is one-for-one in the billed price above $\bar{p}$, unconstrained by any demand response. Government outlays G^{cap} grow as enforcement weakens ($\alpha\gamma \to 0$), because the monopolist's inflation incentive is checked by penalties rather than by lost sales.

3.3 Ad valorem subsidy ($\mathcal{P}_2$)

Under $\mathcal{P}_2$, the government pays a constant share $\tau \in (0, 1)$ of the transaction price. The consumer pays $(1 - \tau)p$, so quantity is $D\bigl((1 - \tau)p\bigr)$ and government outlays are $G^{adv}(p) = \tau\, p\, D\bigl((1 - \tau)p\bigr)$. The firm's profit is $\pi^{adv}(p; \tau) = (p - c)\, D\bigl((1 - \tau)p\bigr)$, which yields the first-order condition

$$D\bigl((1 - \tau)p\bigr) + (p - c)(1 - \tau)\, D'\bigl((1 - \tau)p\bigr) = 0. \tag{6}$$

Let $p^{adv}(\tau)$ denote the solution to (6). The HCP net cost, quantity, expenditure, and government outlay are $p_c^{adv}(\tau) \equiv (1 - \tau)\, p^{adv}(\tau)$, $Q^{adv}(\tau) \equiv D\bigl(p_c^{adv}(\tau)\bigr)$, $E^{adv}(\tau) \equiv p_c^{adv}(\tau)\, Q^{adv}(\tau)$, and $G^{adv}(\tau) \equiv \tau\, p^{adv}(\tau)\, Q^{adv}(\tau)$, respectively. Unlike the price-cap regime, the firm faces a downward-sloping residual demand in its posted price because consumers bear a share $(1 - \tau)$ of each price increase. $\mathcal{P}_2$ restores the elasticity of demand that was suppressed under $\mathcal{P}_1$. Substituting $p_c \equiv (1 - \tau)p$ into (6) yields:

$$D(p_c) + \bigl(p_c - c(1 - \tau)\bigr)\, D'(p_c) = 0. \tag{7}$$

This is the standard monopoly first-order condition (1) with *effective marginal cost* $c(1 - \tau)$

[4] More generally, because $\Phi'(0) = 0$, exceeding the cap is always locally profitable. Even when $\bar{p} \geq p^{no}$, the cap binds globally if enforcement is sufficiently weak. The insulation rent $D(\bar{p})^2/(2\alpha\gamma)$ can exceed the standard monopoly profit loss from operating at $\bar{p}$ instead of p^{no}. In the RHC context, the empirically relevant case is $\bar{p} < p^{no}$, where the cap binds for all enforcement levels.

replacing c. The ad valorem subsidy acts *as if* it reduces the monopolist's cost by the factor $(1-\tau)$ and lets the consumer face the resulting monopoly price directly:

$$p_c^{adv}(\tau) = p^{no}\big(c(1-\tau)\big), \tag{8}$$

where $p^{no}(\tilde{c})$ denotes the monopoly price as a function of marginal cost $\tilde{c}$. We assume the monopoly problem is regular on $\tilde{c} \in [0, c]$: for each $\tilde{c}$, $p^{no}(\tilde{c})$ is the unique finite global maximizer of $(p - \tilde{c})D(p)$, and $\tilde{c} \mapsto p^{no}(\tilde{c})$ is continuous and strictly increasing. We also impose the endpoint condition

$$p^{no}(0) < \bar{p} < p^{no}(c). \tag{9}$$

The upper inequality is the binding-cap condition, while the lower inequality rules out cases in which even the limiting full subsidy cannot reduce the ad valorem consumer price below the cap.[5] This effective-marginal-cost representation (8) characterizes the condition for an HCP to voluntarily switch from $\mathcal{P}_1$ to $\mathcal{P}_2$. By regularity, the consumer price $p_c^{adv}(\tau) = p^{no}\big(c(1-\tau)\big)$ is continuous and strictly decreasing in τ, with $p_c^{adv}(0) = p^{no}(c)$ and $\lim_{\tau\uparrow 1} p_c^{adv}(\tau) = p^{no}(0)$. Therefore, condition (9) guarantees the existence of a unique *critical subsidy rate* $\tau^* \in (0,1)$ satisfying:

$$p_c^{adv}(\tau^*) = \bar{p}, \tag{10}$$

A cost-minimizing HCP switches from $\mathcal{P}_1$ to $\mathcal{P}_2$ whenever $\tau \geq \tau^*$, i.e., whenever the ad valorem consumer price falls below the cap.

Proposition P.2 (Ad valorem dominance)**.** Suppose the regularity condition above and endpoint condition (9) hold, and $\tau \geq \tau^*$. Then:

(i) $p_c^{adv}(\tau) \leq \bar{p} = p_c^{cap}$, with strict inequality when $\tau > \tau^*$;

(ii) $Q^{adv}(\tau) \geq Q^{cap}$;

(iii) if $\varepsilon_D(p) \leq 1$ on $[p_c^{adv}(\tau), \bar{p}]$, then $E^{adv}(\tau) \leq E^{cap}$;

(iv) under the quadratic penalty, $G^{adv}(\tau) < G^{cap}$ for sufficiently small $\alpha\gamma$.

Proof. See Appendix B.2. □

Proposition **P.2** establishes that the ad valorem mechanism dominates the price cap on every margin relevant to policymakers under empirically relevant conditions. The HCP pays less per unit under $\mathcal{P}_2$ (Part i). A lower consumer price raises quantity consumed (Part ii). HCP expenditure falls when demand is price-inelastic (Part iii). Government outlays decline when enforcement is weak (Part iv), because the consumer absorbs a share $(1-\tau)$ of the price, anchoring the monopolist's markup to the demand curve rather than to the enforcement technology.[6]

[5] Under linear demand, $\tau^* = 2(p^{no} - \bar{p})/c$. The condition $\tau^* < 1$ requires $\bar{p} > p^{no} - c/2$; when this fails, even full subsidy coverage cannot offset the monopolist's markup. See Appendix B.4.

[6] Part (iv) is a sufficient-condition existence result: it establishes the existence of an enforcement threshold $\overline{\alpha\gamma}$ below which government outlays under $\mathcal{P}_2$ fall short of those under $\mathcal{P}_1$, but does not pin the threshold down without further information on the penalty function Φ. The empirical exercise in Section 5 tests the implication directly

3.4 Consortium under ad valorem ($\mathcal{P}_2^c$)

Mechanism $\mathcal{P}_2^c$ allows eligible HCPs to form a consortium with ineligible HCPs. Although subsidies apply only to eligible members, the consortium can reallocate billed charges across members while keeping total ISP revenue and consumed quantities fixed. Shifting charges toward eligible members increases government reimbursements, which finance lower prices for ineligible members. Eligible HCPs thus become a conduit for subsidies to flow to ineligible participants. We develop a stylized model of this channel.

The monopolist ISP serves two HCPs. An *eligible* member E receives an ad valorem subsidy at rate $\tau \in (0, 1)$, and an *ineligible* member I has no subsidy. The ISP determines per-unit prices (p_E, p_I), yielding consumer prices $p_{c,E} = (1 - \tau)\, p_E$, $p_{c,I} = p_I$ and quantities $Q_E = D_E(p_{c,E})$ and $Q_I = D_I(p_{c,I})$. The firm chooses prices to maximize profit

$$\pi(p_E, p_I) = (p_E - c)\, D_E\big((1 - \tau)p_E\big) + (p_I - c)\, D_I(p_I).$$

The equilibrium prices (p_E^{adv}, p_I^{adv}) satisfy the first-order conditions in each market separately:

$$\begin{aligned} &D_E\big((1-\tau)p_E^{adv}\big) + (p_E^{adv} - c)(1-\tau)\, D_E'\big((1-\tau)p_E^{adv}\big) \;=\; 0, \\ &D_I(p_I^{adv}) + \big(p_I^{adv} - c\big)\, D_I'(p_I^{adv}) \;=\; 0. \end{aligned}$$

The consortium sets an internal price allocation $(\tilde{p}_E, \tilde{p}_I)$ that aims to reduce total HCP net cost while holding quantity levels and total ISP revenue fixed:

$$p_E^{adv} Q_E^{adv} + p_I^{adv} Q_I^{adv} = \tilde{p}_E\, Q_E^{adv} + \tilde{p}_I\, Q_I^{adv}, \qquad \tilde{p}_E, \tilde{p}_I \geq 0.$$

The reduction in total HCP net cost satisfies the accounting identity

$$\Delta C = \tau\, (\tilde{p}_E - p_E^{adv})\, Q_E^{adv} = \tau\, (p_I^{adv} - \tilde{p}_I)\, Q_I^{adv}. \tag{11}$$

Each dollar shifted to the eligible member's bill reduces the combined HCP net cost by τ dollars, financed by additional government reimbursements. This is a zero-sum transfer between the consortium and the government budget. The government pays $\tilde{G} = \tau \tilde{p}_E\, Q_E^{adv}$ and the increase in the government outlay is $\Delta G = \tau(\tilde{p}_E - p_E^{adv})\, Q_E^{adv} = \Delta C$. We summarize the consortium's problem with three reduced-form parameters:

$$\kappa \equiv \frac{\tilde{p}_E}{p_E^{adv}} \geq 1, \qquad B \equiv \tau\, p_E^{adv} Q_E^{adv}, \qquad R \equiv \frac{p_I^{adv} Q_I^{adv}}{p_E^{adv} Q_E^{adv}} \geq 0. \tag{12}$$

The *eligible price-distortion ratio* κ measures how much the consortium inflates the eligible member's internal price above the equilibrium level. $\kappa = 1$ corresponds to no distortion, while $\kappa > 1$ indicates cross-subsidization. The *eligible subsidy base* B is the government reimbursement at the undistorted equilibrium, i.e., the baseline subsidy before any reallocation. The *ineligible-to-eligible revenue ratio* R measures consortium composition. $R = 0$ means no ineligible members, while large R means ineligible revenue dominates. Nonnegativity of $\tilde{p}_I$ requires $\kappa \in [1, 1 + R]$, since

rather than estimating the regime; the sign of $\tau_{1,2}$ on government outlays therefore serves as the empirical analog of the regime test.

the consortium cannot shift more revenue to the eligible bill than the ineligible member generates. Under this notation, the cost reduction becomes $\Delta C = B(\kappa - 1)$.

Cross-subsidization is constrained by an expected penalty $\alpha\,\Phi(\Delta C;\, R)$, where $\alpha \in (0,1]$ is the probability of being audited and $\Phi : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ is a penalty function satisfying, for every $R > 0$, five conditions: (i) $\Phi(0;\, R) = 0$: no distortion incurs no penalty; (ii) $\partial\Phi/\partial\delta(0;\, R) = 0$: the marginal penalty at zero distortion is zero, ensuring that small deviations are not deterred; (iii) $\partial\Phi/\partial\delta(\delta;\, R) > 0$ for $\delta > 0$: the penalty is strictly increasing in the distortion; (iv) $\partial^2\Phi/\partial\delta^2(\delta;\, R) > 0$: the penalty is strictly convex, so that marginal penalties escalate with the degree of cross-subsidization; and (v) $\partial^2\Phi/\partial\delta\,\partial R > 0$: the penalty is supermodular in (δ, R), so that consortia with larger ineligible-to-eligible revenue ratios face a steeper marginal deterrent. For tractability, we adopt the quadratic specification $\Phi(\delta;\, R) = \frac{\gamma R}{2}\,\delta^2$ $(\gamma > 0)$, which satisfies these conditions on the nondegenerate domain $R > 0$ and yields closed-form solutions. The consortium's price allocation solves

$$\min_{\tilde{p}_E,\, \tilde{p}_I \geq 0} \quad \underbrace{(1-\tau)\,\tilde{p}_E\, Q_E^{adv} \;+\; \tilde{p}_I\, Q_I^{adv}}_{\text{HCP net cost}} \;+\; \underbrace{\alpha\,\Phi(\Delta C;\, R)}_{\text{expected penalty}} \tag{13}$$

$$\text{s.t.} \quad \tilde{p}_E\, Q_E^{adv} + \tilde{p}_I\, Q_I^{adv} \;=\; p_E^{adv}\, Q_E^{adv} + p_I^{adv}\, Q_I^{adv},$$

which can be rewritten as a maximization of cost reduction minus the expected penalty over the price-distortion ratio κ:

$$\max_{\kappa \in [1,\, 1+R]} \; \Psi(\kappa), \qquad \Psi(\kappa) \equiv B\,\kappa - \alpha\,\Phi\big(B(\kappa-1);\, R\big). \tag{14}$$

Proposition P.3 (Optimal cross-subsidization)**.** Suppose $B > 0$, $R > 0$, and Φ satisfies conditions *(i)–(v)*. Then:

(i) Cross-subsidization is locally profitable at no distortion: $\Psi'(1) = B > 0$, which implies $\kappa^* > 1$ and $\tilde{p}_E > p_E^{adv}$.

(ii) Ψ is strictly concave on $[1,\, 1+R]$. Hence the optimal distortion κ^* is unique.

(iii) If $\alpha\,\partial\Phi/\partial\delta(BR;\, R) > 1$, the optimum is interior $(1 < \kappa^* < 1+R)$ and satisfies
$$\alpha\,\partial\Phi/\partial\delta\big(B(\kappa^*-1);\, R\big) = 1. \tag{15}$$
If $\alpha\,\partial\Phi/\partial\delta(BR;\, R) \leq 1$, the constraint binds and $\kappa^* = 1 + R$.

(iv) Under the quadratic penalty $\Phi(\delta;\, R) = \frac{\gamma R}{2}\,\delta^2$,
$$\kappa^*(R) = \min\left\{1+R,\; 1 + \frac{1}{\alpha\gamma B\, R}\right\}. \tag{16}$$
The optimal distortion κ^* is hump-shaped in R, attaining its maximum $\kappa^* = 1 + R^*$ at $R^* = 1\big/\sqrt{\alpha\gamma B}$, and converging to 1 (no cross-subsidization) as $R \downarrow 0$ or $R \to \infty$. At the boundary $R = 0$, feasibility forces $\kappa^* = 1$.

Proof. See Appendix B.3. □

Part (i) establishes that some cross-subsidization occurs whenever there is positive ineligible revenue to reallocate. Because $\partial\Phi/\partial\delta(0; R) = 0$, the marginal penalty at zero distortion is zero while the marginal saving is $B > 0$. If $R = 0$, the feasibility set collapses to $\kappa = 1$, so no cross-subsidization is possible. Part (ii) guarantees uniqueness via the convexity of the penalty. Part (iii) characterizes the two regimes. When R is small but positive, there is little ineligible revenue to shift, so the feasibility constraint $\kappa \le 1 + R$ binds before enforcement does. When R is large, the marginal penalty at full reallocation exceeds the marginal saving, so enforcement binds first and the consortium self-restrains. Part (iv) shows that under the quadratic penalty the transition between the two regimes occurs at $R^* = 1/\sqrt{\alpha\gamma B}$, generating a hump shape: $\kappa^*(R)$ traces the tension between cross-subsidization opportunity and enforcement exposure, analogous to the Laffer curve (see Figure 5).

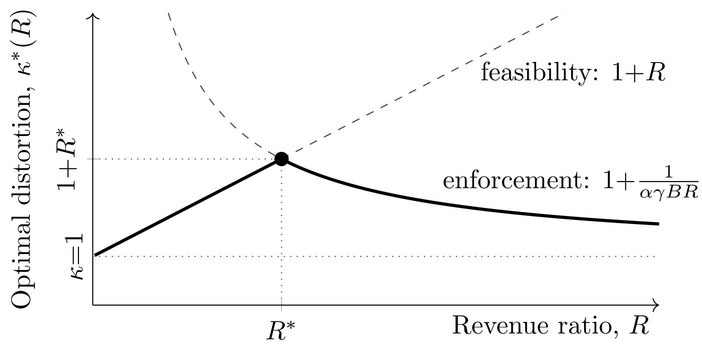

Figure 5: Optimal cross-subsidization distortion as a function of consortium composition.

Notes: Solid curve: $\kappa^*(R) = \min\{1 + R,\ 1 + 1/(\alpha\gamma BR)\}$ (Proposition **P.3**, Part iv). The peak at $R^* = 1/\sqrt{\alpha\gamma B}$ maximizes the distortion, a Laffer-curve logic applied to subsidy extraction. For small R, feasibility binds (insufficient ineligible revenue); for large R, enforcement binds. Dashed lines show each branch beyond its binding region. Illustration uses $\alpha\gamma B = 1$.

3.5 Empirical implications

The three mechanisms generate testable predictions. The 2014 introduction of HCF made $\mathcal{P}_2$ and $\mathcal{P}_2^c$ available to incumbent $\mathcal{P}_1$ HCPs for the first time, and identification therefore exploits the two observable transitions $\mathcal{P}_1$ to $\mathcal{P}_2$ and $\mathcal{P}_1$ to $\mathcal{P}_2^c$.

Proposition **P.2** predicts that switching from the price cap $\mathcal{P}_1$ to the ad valorem $\mathcal{P}_2$ lowers the consumer price (Part i), raises quantity (Part ii), reduces HCP expenditure when demand is inelastic (Part iii), and lowers government outlays when enforcement is weak (Part iv). The key economic force is that the ad valorem subsidy restores the link between the billed price and demand, eliminating the insulation rent that arises under the price cap.

An HCP switching from $\mathcal{P}_1$ to $\mathcal{P}_2^c$ experiences both forces simultaneously: the efficiency gain from ad valorem pricing and the distortion from consortium reallocation (Proposition **P.3**). Because the consortium inflates the eligible member's billed price above the ad valorem equilibrium ($\tilde{p}_E > p_E^{adv}$), the cost reduction from switching to $\mathcal{P}_2^c$ should be smaller than from switching to $\mathcal{P}_2$. When the distortion dominates the efficiency gain, the transition can increase costs. The difference $\tau_{1,2^c} - \tau_{1,2}$ provides a model-consistent measure of the cross-subsidization premium, assuming the two switcher populations have comparable selection on unobservables.

Finally, Proposition **P.3**, Part (iv), predicts that the price distortion $\kappa^*(R)$ first rises and then falls with the ineligible-to-eligible revenue ratio R, holding the eligible-side bargaining base $B \equiv \tau\, p_E^{adv} Q_E^{adv}$ fixed. This is a *within-*$\mathcal{P}_2^c$ ceteris-paribus prediction: among consortia with comparable eligible-side fundamentals, the relationship between price inflation and the ineligible share should exhibit an inverted U. In the empirical implementation, cross-consortium variation in R co-varies with B through primitives such as τ and demand parameters, so the empirical test in Section 5.8 approximates the ceteris-paribus shape by partialling out HCP fixed effects and observable controls; it should be read as a model-consistent shape test rather than a structural estimate. These implications motivate two hypotheses that structure the empirical analysis.

Hypothesis H.1 (Cost effectiveness)**.** The ad valorem mechanism $\mathcal{P}_2$ reduces prices and government outlays relative to the price cap $\mathcal{P}_1$.

Hypothesis H.2 (Cross-subsidization)**.** Cross-subsidization under $\mathcal{P}_2^c$ erodes the ad valorem efficiency gain. The cost reduction from switching to $\mathcal{P}_2^c$ is smaller than from switching to $\mathcal{P}_2$, and when the distortion is sufficiently large it can offset the intended benefits entirely.

Together, these hypotheses capture the paper's central tension. The 2014 reform corrected a design flaw in subsidy delivery, but inadvertently opened a channel for strategic manipulation that partially undoes the correction. Section 4 develops the econometric framework, and Section 5 tests both hypotheses.

4 Econometric framework

4.1 Setup and identification

Our empirical setting is a two-period DiD model with two treatment margins. Let $Y_{it}(d)$ denote the potential outcome for HCP i at time $t \in \{0, 1\}$ under treatment status $d \in \{1,\ 2,\ 2^c\}$, where $d = 1$ corresponds to $\mathcal{P}_1$, $d = 2$ to $\mathcal{P}_2$, and $d = 2^c$ to $\mathcal{P}_2^c$. Period $t = 0$ is 2013 and $t = 1$ is 2014. Because $\mathcal{P}_2$ and $\mathcal{P}_2^c$ were introduced in 2014, all HCPs are on $\mathcal{P}_1$ at $t = 0$. In 2014, each HCP either remains on $\mathcal{P}_1$ or switches to $\mathcal{P}_2$ or $\mathcal{P}_2^c$.

This yields three treatment groups. Let $g_{d,d'}$ denote the group of HCPs in status d at $t = 0$ and status d' at $t = 1$. Since all HCPs start on $\mathcal{P}_1$, the set of groups is $\mathcal{G} \equiv \{g_{1,1},\ g_{1,2},\ g_{1,2^c}\}$, where $g_{1,1}$ denotes stayers, $g_{1,2}$ denotes those that switch to $\mathcal{P}_2$, and $g_{1,2^c}$ denotes those that switch to $\mathcal{P}_2^c$. We write $D_{i,2} \equiv \mathbf{1}[G_i = g_{1,2}]$ and $D_{i,2^c} \equiv \mathbf{1}[G_i = g_{1,2^c}]$ for the treatment group indicators, and $T_t \equiv \mathbf{1}[t = 1]$ for the post-treatment period.

We target two average treatment effects on the treated (ATTs), one for each switching margin. Table 1 lists each ATT alongside its identifying parallel trends assumption. The control group consists of HCPs that remain on $\mathcal{P}_1$ in both periods ($g_{1,1}$). The unit is the individual eligible HCP, and within-consortium reallocation is internalized in the $\mathcal{P}_2^c$ treatment definition. The stable unit treatment value assumption (SUTVA) then requires only that distinct consortia do not affect each other (they bill USAC independently) and that switchers do not affect $\mathcal{P}_1$ stayers (per-unit prices are negotiated bilaterally).[7]

Table 1—Group-specific treatment effects and parallel trends assumptions.

Treatment Effect (ATT)	Parallel Trends Assumption
$\tau_{1,2} \equiv \mathbb{E}[Y_{i1}(2) - Y_{i1}(1) \mid G_i = g_{1,2}]$	$\mathbf{PT}_{1,2}$: $\mathbb{E}[Y_{i1}(1) - Y_{i0}(1) \mid G_i = g_{1,2}] = \mathbb{E}[Y_{i1}(1) - Y_{i0}(1) \mid G_i = g_{1,1}]$
$\tau_{1,2^c} \equiv \mathbb{E}[Y_{i1}(2^c) - Y_{i1}(1) \mid G_i = g_{1,2^c}]$	$\mathbf{PT}_{1,2^c}$: $\mathbb{E}[Y_{i1}(1) - Y_{i0}(1) \mid G_i = g_{1,2^c}] = \mathbb{E}[Y_{i1}(1) - Y_{i0}(1) \mid G_i = g_{1,1}]$

Notes: $Y_{it}(d)$ is the potential outcome of HCP i in period $t \in \{0,1\}$ under program $d \in \{1,2,2^c\}$. G_i is the group assignment: $g_{1,2}$ switchers from $\mathcal{P}_1$ to $\mathcal{P}_2$, $g_{1,2^c}$ switchers from $\mathcal{P}_1$ to $\mathcal{P}_2^c$, and $g_{1,1}$ stayers on $\mathcal{P}_1$.

Under parallel trends, the ATTs are point identified as:

$$\tau_{1,2} = \mathbb{E}\big[Y_{i1} - Y_{i0} \mid G_i = g_{1,2}\big] - \mathbb{E}\big[Y_{i1} - Y_{i0} \mid G_i = g_{1,1}\big], \tag{17}$$

$$\tau_{1,2^c} = \mathbb{E}\big[Y_{i1} - Y_{i0} \mid G_i = g_{1,2^c}\big] - \mathbb{E}\big[Y_{i1} - Y_{i0} \mid G_i = g_{1,1}\big]. \tag{18}$$

The theoretical model in Section 3 yields predictions for these parameters. The cost-effectiveness channel implies $\tau_{1,2} < 0$: switching from the price cap to the ad valorem mechanism should lower prices because the HCP now bears a share of the cost and resists price inflation. The parameter $\tau_{1,2^c}$ captures cost-effectiveness and cross-subsidization simultaneously, since switching from $\mathcal{P}_1$ to $\mathcal{P}_2^c$ changes both the reimbursement formula and the group structure. The indirect test $\tau_{1,2^c} - \tau_{1,2} > 0$ therefore isolates the cross-subsidization component. A positive difference indicates that consortium membership inflates eligible members' costs beyond what the ad valorem mechanism alone would produce.

4.2 Estimation

Since all HCPs share the same baseline status ($\mathcal{P}_1$ at $t = 0$), the estimating equation is the standard multi-treatment TWFE:

$$Y_{it} \;=\; \alpha \;+\; \gamma_2\, D_{i,2} \;+\; \gamma_{2^c}\, D_{i,2^c} \;+\; \lambda\, T_t \;+\; \tau_{1,2}\big(D_{i,2} \times T_t\big) \;+\; \tau_{1,2^c}\big(D_{i,2^c} \times T_t\big) \;+\; \varepsilon_{it}, \tag{19}$$

where α is the intercept, γ_2 and γ_{2^c} are group fixed effects, λ is the period effect, and ε_{it} is a mean-zero error term. In the binary specification, the coefficients $\tau_{1,2}$ and $\tau_{1,2^c}$ recover the ATTs in equations (17)–(18).[8]

4.3 Continuous program-share specification

Each HCP may comprise multiple facilities, such as a hospital chain with several campuses, each of which independently chooses among $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_2^c$. Because the data lack facility identifiers but map every facility to its parent HCP, we work at the HCP level. Let $\mathcal{I}_j$ denote the set of facilities

[7] The remaining standard assumptions are no anticipation and common support $0 < \Pr(G_i = g) < 1$ for all $g \in \mathcal{G}$.

[8] In a balanced panel, equation (19) simplifies to first differences: $\Delta Y_{it} = \lambda + \tau_{1,2}\, D_{i,2} + \tau_{1,2^c}\, D_{i,2^c} + \Delta\varepsilon_{it}$, where $\Delta Y_{it} \equiv Y_{i1} - Y_{i0}$.

in HCP j, with quantity (Mbps) Q_{it} and price Y_{it} observed at the facility level. We construct the quantity-weighted average price for HCP j in period t:

$$\bar{Y}_{jt} = \frac{\sum_{i \in \mathcal{I}_j} Y_{it} \, Q_{it}}{\sum_{k \in \mathcal{I}_j} Q_{kt}},$$

and define the program-share variables

$$S_{jt,2} = \frac{\sum_{i \in \mathcal{I}_j} D_{i,2} \, Q_{it}}{\sum_{k \in \mathcal{I}_j} Q_{kt}}, \qquad S_{jt,2^c} = \frac{\sum_{i \in \mathcal{I}_j} D_{i,2^c} \, Q_{it}}{\sum_{k \in \mathcal{I}_j} Q_{kt}},$$

as the fractions of HCP j's bandwidth allocated to $\mathcal{P}_2$ and $\mathcal{P}_2^c$ at time t, respectively. These shares replace the binary indicators in equation (19):

$$\bar{Y}_{jt} \;=\; \alpha \;+\; \gamma_2 \, S_{jt,2} \;+\; \gamma_{2^c} \, S_{jt,2^c} \;+\; \lambda \, T_t \;+\; \tau_{1,2}\big(S_{jt,2} \times T_t\big) \;+\; \tau_{1,2^c}\big(S_{jt,2^c} \times T_t\big) \;+\; \xi_{jt}, \tag{20}$$

where ξ_{jt} is the quantity-weighted average of facility-level errors.

The coefficients $\tau_{1,2}$ and $\tau_{1,2^c}$ in equation (20) recover the average treatment effect on the treateds (ATTs) in (17)–(18) under parallel trends and an aggregation restriction: the quantity-weighted mean facility-level effect of switching from $\mathcal{P}_1$ to destination program d is invariant to the HCP-level program-share vector $(S_{jt,2}, S_{jt,2^c})$, after conditioning on any included controls. This restriction allows treatment effects to vary across facilities, but rules out systematic correlation between treatment-effect heterogeneity and the share variation used to estimate the coefficients. Intermediate values of $S_{jt,2}$ and $S_{jt,2^c}$ therefore scale the underlying discrete transition effects as quantity-weighted aggregates rather than defining a separate continuous treatment. Appendix C provides the formal identification argument and characterizes the population coefficients as activity-weighted ATT aggregates under this aggregation restriction.[9]

4.4 Covariates and inference

We augment equation (20) with a vector of observed covariates $\boldsymbol{X}_{jt}$ comprising log internet speed, HCP type indicators, internet service type indicators, state fixed effects, and year fixed effects. The linear specification is:

$$\bar{Y}_{jt} \;=\; \boldsymbol{S}_{jt}' \, \boldsymbol{\theta}_0 \;+\; \boldsymbol{X}_{jt}' \, \boldsymbol{\beta}_0 \;+\; \xi_{jt}, \tag{21}$$

where $\boldsymbol{S}_{jt}$ stacks the treatment share variables and their time interactions, $\boldsymbol{\theta}_0$ collects the corresponding parameters including the ATTs of interest, and $\boldsymbol{\beta}_0$ captures the covariate effects. Standard errors are clustered at the HCP level throughout.

Linear control for $\boldsymbol{X}_{jt}$ may be misspecified if the relationship between covariates and outcomes is nonlinear. To address this, we supplement the baseline with a DML estimator (Chernozhukov et al., 2018), which replaces the linear control function with a flexible, unknown function $g_0(\boldsymbol{X}_{jt})$

[9] Linear TWFE estimators can deliver weighted averages with negative weights when treatment effects are heterogeneous and adoption is staggered (Borusyak, Jaravel, and Spiess, 2024; de Chaisemartin and d'Haultfœuille, 2020), and dedicated estimators for scalar continuous-dose DiD have been developed (Callaway, Goodman-Bacon, and Sant'Anna, 2024). Our setting differs because the primitive treatment choices are discrete and the observed shares are HCP-level aggregates. Since all units share a single pre-period (2013) and a single post-period (2014), adoption is not staggered; the maintained aggregation restriction above is what gives the share coefficients their ATT interpretation.

estimated via random forest. Under the high-level rate and regularity conditions for the nuisance functions, Neyman-orthogonal moments and K-fold cross-fitting (K=10) deliver $\sqrt{n}$-consistent inference on $\boldsymbol{\theta}_0$ while reducing first-order sensitivity to regularization bias. Appendix D details the implementation.

5 An empirical exercise

5.1 Data

We use the RHC Commitments and Disbursements data provided by USAC. The data is published under an FCC mandate.[10] It covers the full population of HCPs who sought $\mathcal{P}_1$, $\mathcal{P}_2$, or $\mathcal{P}_2^c$ subsidies, whether approved, denied, or pending. The panel spans 2012 to the present, with new records added on a rolling basis. Our analysis uses observations through 2021. Three identifiers structure the panel: funding year, HCP ID, and consortium ID.

The data is at the "line of subsidy" level. A subsidy request may include one or multiple lines of internet connection. Consider first a single-line request. It is recorded as one observation, with a unique Funding Request Number (FRN) identifying the request. Because the request has only one line of subsidy, its FRN Line Number (FRNLN) equals 1. If a request includes multiple lines of internet connection, each line is recorded as a separate observation. All lines share the same FRN. Each is identified by an FRNLN that starts at 1 and increments by 1. In consortium applications, all lines also share a consortium ID, and each request maps to its HCP via the HCP ID.

The data underwent an extensive clean-up and exclusion process, detailed in Appendix E and summarized in Table A3. The 2013–2014 sample starts with 36,576 request-level observations. We limit the sample to annual contracts and to categories of expense that correspond to internet subscription costs (excluding equipment, maintenance, and infrastructure). We remove grandfathered requests, pending and withdrawn requests, and cases missing speed, price, or subsidy. We drop Alaska, an outlier state with exorbitant prices, the lowest speeds, and a unique technology profile. After limiting speeds to 1–100megabit per second (Mbps) and dropping zero-subsidy observations, 7,780 request-level observations remain.

To enable panel analysis, we aggregate observations to the HCP–year level. Each HCP's annual speed is the sum of speeds across its lines in that year. Price and subsidy are aggregated the same way. We use the number of lines of subsidy per HCP as a measure of HCP size. After aggregation, the 2013–2014 sample contains 4,231 HCP-year observations. Restricting to HCPs that appear in both years produces a balanced panel of 1,940 observations (970 HCPs observed twice). The same clean-up procedure is applied to each subsequent year pair through 2020–2021, yielding balanced panels that grow to 7,812 HCP-year observations in 2020–2021 (Table A9). These year pairs feed the combined estimates in Section 5.7.

[10] https://opendata.usac.org/Rural-Health-Care/RHC-Commitments-and-Disbursements-Tool/sm8n-gg82

Table 2—Program utilization, 2013–2014.

	2013	2014		
	$\mathcal{P}_1$	$\mathcal{P}_1$	$\mathcal{P}_2$	$\mathcal{P}_2^c$
Number of HCPs	970	643	419	43
Number of requests	1,938	1,263	850	8
Number of lines	1,980	1,305	973	56
Percentage of lines	100.0%	55.9%	41.7%	2.4%
Percentage of speed	100.0%	35.3%	63.1%	1.7%
Percentage of subsidies	100.0%	77.9%	16.5%	5.5%

Notes: Sample of 970 baseline HCPs. A request may include multiple lines. Percentage rows give each program's share of the year's total lines, speed, and subsidy outlays.

Table 2 reports approved subsidy requests, HCPs, and lines of internet connection by program for the baseline 2013–2014 sample. By 2014, 419 of the 970 incumbents had switched to $\mathcal{P}_2$ and 43 to $\mathcal{P}_2^c$, while 643 remained on $\mathcal{P}_1$. These counts sum to more than 970 because an HCP drawing from multiple programs in a year is counted under each. Although $\mathcal{P}_2^c$ accounts for only 2.4% of lines, it absorbs 5.5% of subsidies, an early indicator that consortium filing carries a higher per-line cost.

The three programs share the same eligibility criteria, coverage, and bureaucracy. Rational HCPs should therefore choose the program that minimizes HCP net cost. Yet in 2014, HCPs who switched from $\mathcal{P}_1$ to $\mathcal{P}_2^c$ accepted unusually high net costs. Cross-subsidization rationalizes the anomaly: an HCP will accept a higher net cost if the resulting cross-subsidization gains exceed it. Because 2014 was the first year consortium filing was available, the HCPs with the largest expected gains had the strongest incentive to switch immediately. This first-mover pattern shows up as the early jump in HCP net cost.

5.2 A descriptive regression analysis

Table A7 reports summary statistics at the subsidy-line level. Price, subsidy, and HCP net cost vary widely. In the 2013–2014 sample, $\mathcal{P}_1$ accounts for 75.7% of observations (3,205 of 4,234), with $\mathcal{P}_2$ at 23.0% and $\mathcal{P}_2^c$ at 1.3%. The sample is dominated by non-profit hospitals, rural health clinics, and community health centers, with Ethernet and internet as the leading service types. Observations span most US states, concentrated in California, Wisconsin, Texas, and Illinois. We estimate the following specification:

$$\ln(LHS_{it}) = \beta_1\mathcal{P}_1 + \beta_2\mathcal{P}_2 + \beta_3\mathcal{P}_2^c + \gamma \ln \mathit{speed} + \lambda_t + \alpha_s + \psi_i + \theta_z + \epsilon_{it}$$

where LHS_{it} is one of price, subsidy, or HCP net cost. The terms λ_t, α_s, ψ_i, θ_z are fixed effects (FE) for year, state, HCP type, and internet service type. The error term is ϵ_{it}.

Table 3 reports the estimation results. Internet plans are cheaper under $\mathcal{P}_2$ than $\mathcal{P}_1$ (Column 2). This is consistent with the cost-efficiency hypothesis. $\mathcal{P}_2$'s design restores HCPs' price sensitivity and lowers negotiated prices. The same pattern holds for subsidy and HCP net cost (Columns 3–4).

TABLE 3—POOLED OLS REGRESSION RESULTS.

	ln(price)	ln(subsidy)	ln(HCP net cost)
β_1	7.1634	7.0166	5.3150
	(0.1345)	(0.1759)	(0.2261)
β_2	5.6567	5.5327	4.4648
	(0.1363)	(0.1782)	(0.2290)
β_3	7.2355	6.8700	6.3324
	(0.1668)	(0.2181)	(0.2802)
$\beta_3 - \beta_2$	1.5787	1.3373	1.8675
	(0.1092)	(0.1427)	(0.1835)
ln(speed)	0.3159	0.2827	0.3059
	(0.0168)	(0.0220)	(0.0282)
N	4,234	4,234	4,234
Adj. R^2	0.9946	0.9898	0.9789

Notes: Sample of 970 baseline HCPs over 2013–2014. β_1, β_2, β_3 are coefficients on $\mathcal{P}_1$, $\mathcal{P}_2$, $\mathcal{P}_2^c$ indicators in a no-intercept specification with year, HCP-type, service-type, and state fixed effects. Observations with non-positive HCP net cost are dropped. Standard errors in parentheses.

$\mathcal{P}_2$ reduces both relative to $\mathcal{P}_1$, consistent with cost savings being shared between the government and the HCPs. The high adjusted R^2 indicates that the covariates explain most of the outcome variation.

Because $\mathcal{P}_2$ and $\mathcal{P}_2^c$ share the same cost-sharing rate, any systematic $\mathcal{P}_2^c$–$\mathcal{P}_2$ price gap is suggestive evidence of cross-subsidization. The estimates for $\beta_3 - \beta_2$ are large and highly significant across all three outcomes. The difference is largest for HCP net cost, indicating that consortium members bear a disproportionate share of the price inflation. Appendix F.1 (Table A5) reports a variant that replaces the contemporaneous ln(speed) control with each HCP's pre-period (2013) log-bandwidth. The cross-subsidization premium $\beta_3 - \beta_2$ remains highly significant across all three outcomes, with magnitudes about 26–30% smaller than under contemporaneous control.

5.3 Main results

We estimate the continuous-share model (eq. 20), its binary counterpart (eq. 19), and a DML specification on the baseline 2013–2014 panel. Table 4 reports the results. The estimates are negative across all three outcomes and all three methods, supporting the cost-effectiveness prediction $\tau_{1,2} < 0$. Switching from $\mathcal{P}_1$ to $\mathcal{P}_2$ reduces prices, subsidies, and HCP net cost.

The direct cross-subsidization effect $\tau_{2,2^c}$ is not estimable because $\mathcal{P}_2$ did not exist in 2013. As explained in Section 4, the contrast $\tau_{1,2^c} - \tau_{1,2}$ isolates the cross-subsidization component, and it is positive and significant in Table 4, consistent with the cross-subsidization hypothesis. The binary treatment specification (Panel B) discards HCPs with mixed program participation within a year, yielding a slightly smaller sample. The DML estimates (Panel C) allow covariates to enter the model nonlinearly. Both alternatives are close to the TWFE baseline and highly significant, suggesting the cross-subsidization result is not an artifact of the linear specification or the choice

Table 4—Combined regression results, 2014.

	Panel A: ln(**price**)			Panel B: ln(**subsidy**)			Panel C: ln(**HCP net cost**)		
	Cont.	Binary	DML	Cont.	Binary	DML	Cont.	Binary	DML
$\tau_{1,2}$	−1.249	−1.326	−1.755	−1.262	−1.349	−1.697	−0.704	−0.731	−1.295
	(0.085)	(0.097)	(0.212)	(0.102)	(0.116)	(0.225)	(0.074)	(0.082)	(0.207)
$\tau_{1,2^c}$	0.556	0.600	0.569	0.284	0.310	0.381	1.740	1.863	1.431
	(0.136)	(0.158)	(0.264)	(0.164)	(0.190)	(0.269)	(0.119)	(0.135)	(0.262)
$\tau_{1,2^c} - \tau_{1,2}$	1.804	1.926	2.324	1.546	1.659	2.078	2.445	2.594	2.726
	(0.147)	(0.171)	(0.372)	(0.177)	(0.206)	(0.375)	(0.128)	(0.146)	(0.398)
N	1,940	1,670	1,940	1,940	1,670	1,940	1,940	1,670	1,940
R^2	0.387	0.390		0.312	0.311		0.431	0.448	

Notes: Sample is the 970 baseline HCPs observed in 2013 and 2014 (1,940 observations; 1,670 in the binary specification, which restricts to switchers and stayers). $\tau_{1,2}$ and $\tau_{1,2^c}$ are the effects of switching from $\mathcal{P}_1$ to $\mathcal{P}_2$ and from $\mathcal{P}_1$ to $\mathcal{P}_2^c$. Cont.: TWFE with continuous treatment shares; Binary: TWFE with binary indicators; DML: double machine learning with random forests (10-fold cross-validation). The standard error of $\tau_{1,2^c} - \tau_{1,2}$ under DML is computed from the joint influence-function covariance matrix. A levels analog appears in Table A12. Standard errors in parentheses.

of estimator.

Appendix F reports eight additional robustness checks varying the sample composition (medical-school exclusion, specific HCP types, Alaska included) and data-processing assumptions (all speed ranges, all contract durations, Mbps-only speeds). The results are robust throughout. Appendix F.2 (Table A6) re-estimates the 2013–2014 TWFE/DML table without the ln(speed) control, the cleanest analog of the pre-period-speed robustness used for the Pooled Ordinary Least Squares (POLS). The $\tau_{1,2^c} - \tau_{1,2}$ cross-subsidization premium remains highly significant across all three outcomes and all three estimators, with magnitudes 13–31% smaller than the baseline. The $\tau_{1,2}$ magnitudes shrink by 41–64%, reflecting that the no-ln(speed) specification captures the *total* mechanism effect (including the quantity-channel response) rather than the per-unit-speed effect.

Three appendix checks confirm the baseline. Appendix H applies Cook's Distance diagnostics and re-estimates the model after dropping all influential HCPs. The treatment effects remain significant and economically large. A related concern is that $\mathcal{P}_2^c$ participants operate in a narrower speed range than $\mathcal{P}_1$ or $\mathcal{P}_2$ participants. Appendix I addresses this by imposing a common-support restriction on the $\mathcal{P}_2^c$ speed domain. The coefficients are virtually unchanged. Appendix J applies the Oster (2019) coefficient-stability test. Every δ is negative, so adding time-varying controls moves the treatment coefficients *further* from zero. Unobservables would have to work in the opposite direction of observables to nullify the effects.

Table A10 compares eight price–speed functional forms. The top goodness-of-fit choices are log-log and lin-log. A Box-Cox test (Table A11) rejects both nested specifications, with the maximum-likelihood $\hat{\lambda}$ closer to zero than to one, so we adopt log-log as the baseline. Appendix G (Table A12) reports the lin-log specification as a robustness check, with all key coefficients remaining significant and stable and the implied semi-elasticities of similar magnitude (Table A13). The two specifications bracket the Box-Cox optimum and yield the same conclusions.

5.4 A 2012 placebo

Identification rests on the parallel-trends assumption. Absent treatment, outcomes for switching HCPs would have evolved like those of non-switchers. We test it directly using 2012 data. Because HCF was not yet available before 2013, the 2012→2013 panel is genuinely pre-treatment for all HCPs. We construct a balanced panel of 642 HCPs that appear in all three years (2012, 2013, 2014) and whose 2014 program assignment is identified from the baseline $\mathcal{C}_1 \cap \mathcal{C}_2$ subset. Each HCP's 2013→2014 transition shares $(P_{01}, P_{02}, \ldots)$ are attached as forward-looking transition shares, and we estimate the same continuous-share specification used in the "Cont." columns of Table 4 on the 2012→2013 placebo window. Under parallel trends, the placebo coefficients should be indistinguishable from zero, since future mechanism choice cannot affect 2012→2013 outcomes.

Table 5—Placebo test of parallel trends, 2012 pre-period.

	Panel A: ln(**price**)		**Panel B:** ln(**subsidy**)		**Panel C:** ln(**HCP net cost**)	
	Placebo (2012–13)	Headline (2013–14)	Placebo (2012–13)	Headline (2013–14)	Placebo (2012–13)	Headline (2013–14)
$\tau_{1,2}$	−0.032	−1.452	0.000	−1.523	−0.140	−0.736
	(0.028)	(0.105)	(0.041)	(0.127)	(0.040)	(0.085)
$\tau_{1,2^c}$	0.080	0.580	0.157	0.303	−0.149	1.744
	(0.058)	(0.148)	(0.084)	(0.179)	(0.081)	(0.119)
$\tau_{1,2^c} - \tau_{1,2}$	0.113	2.032	0.157	1.826	−0.009	2.480
	(0.059)	(0.163)	(0.086)	(0.197)	(0.083)	(0.132)
N (HCP-year obs.)	1,284	1,284	1,284	1,284	1,284	1,284
R^2	0.694	0.438	0.580	0.358	0.483	0.514

Notes: The sample is the balanced panel of 642 HCPs appearing in 2012, 2013, and 2014, with 2014 program assignment recovered from the baseline $\mathcal{C}_1 \cap \mathcal{C}_2$ subset. HCF ($\mathcal{P}_2$ and $\mathcal{P}_2^c$) was not available before 2014, so both 2012 and 2013 are pre-treatment. *Headline* columns estimate the baseline specification on the 2013–2014 panel; *Placebo* columns repeat it on the 2012–2013 panel, applying each HCP's 2013–2014 continuous treatment shares as if they had switched a year earlier. Specification matches the "Cont." columns of Table 4: HCP fixed effects with ln(speed) and ln(requests) as controls. Under parallel trends, *Placebo* coefficients should be statistically indistinguishable from zero. Standard errors in parentheses.

Table 5 reports the placebo alongside the 2013→2014 estimates on the same balanced sample. The actual estimates exceed the placebo almost everywhere. In seven of nine cells the placebo is at most 14% of the corresponding 2013→2014 estimate. The cross-subsidization differential $\tau_{1,2^c} - \tau_{1,2}$ that anchors the $\mathcal{P}_2^c$ claim has placebo magnitudes of only 5.6%, 8.6%, and 0.4% of the actual effect for price, subsidy, and HCP net cost. Two cells show larger pre-trends, $\mathcal{P}_2$-arm HCP net cost at 19% of its actual effect and $\mathcal{P}_2^c$-arm subsidy at 52%, but neither overturns the sign or the qualitative conclusion. A third placebo, $\mathcal{P}_2^c$-arm HCP net cost, runs opposite to its actual effect and so biases against our result. The placebo supports a causal reading of the 2013→2014 estimates.

Figure A4 plots the same six estimates as a three-period event study anchored at 2012, separately for the $\mathcal{P}_2$ and $\mathcal{P}_2^c$ arms. The trends are largely flat through 2013 and jump in 2014, with the two pre-period exceptions noted above visible as 2013 deviations.

5.5 Aryal–Manski sensitivity

As a complementary check, we follow Aryal et al. (2025) (building on Manski and Pepper (2018)) in measuring the sensitivity of the results to a parameterized violation of the parallel-trends assumption. Aryal et al. (2025) let the unobserved counterfactual trend for the treatment group (β^1) be proportional to the observed trend of the control group (β^0) by a scalar $g \in [0, 2]$, so $\beta^1 = g \cdot \beta^0$ (standard parallel trends corresponds to $g = 1$). A robust treatment effect, $\hat{\beta}^{\text{Robust}}$, then adjusts for this potential violation:

$$\hat{\beta}^{\text{Robust}} = \hat{\beta}^{\text{DiD}} + \left(\hat{\beta}^0 - \hat{\beta}^1\right) = \hat{\beta}^{\text{DiD}} + (1 - g) \cdot \hat{\beta}^0$$

Applying this test to our specification requires a modification. The Aryal–Manski approach needs the post-period indicator T_t as a stand-alone regressor, which our continuous-share specification does not provide. We restore the canonical form by partitioning the 2014 sample. In 2013 all subsidy requests were in $\mathcal{P}_1$. In 2014 they were split across $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_2^c$. Discarding 2014's $\mathcal{P}_2^c$ observations restricts the remaining HCPs to D_{00} or D_{01}, so $D_{02} = D_{11} = D_{12} = D_{22} = 0$ and eq. 19 reduces to a canonical DiD with a single treatment effect, enabling the test for $\tau_{1,2}$. Symmetrically, discarding 2014's $\mathcal{P}_2$ observations enables the test for $\tau_{1,2^c}$.

Restricting to a subset of the data limits inference, but a result that survives the test under multiple values of g supports robustness of the baseline to mild parallel-trends violations. We run the test for $\tau_{1,2}$ and $\tau_{1,2^c}$, separately for price, subsidy, and HCP net cost, generating six tests.

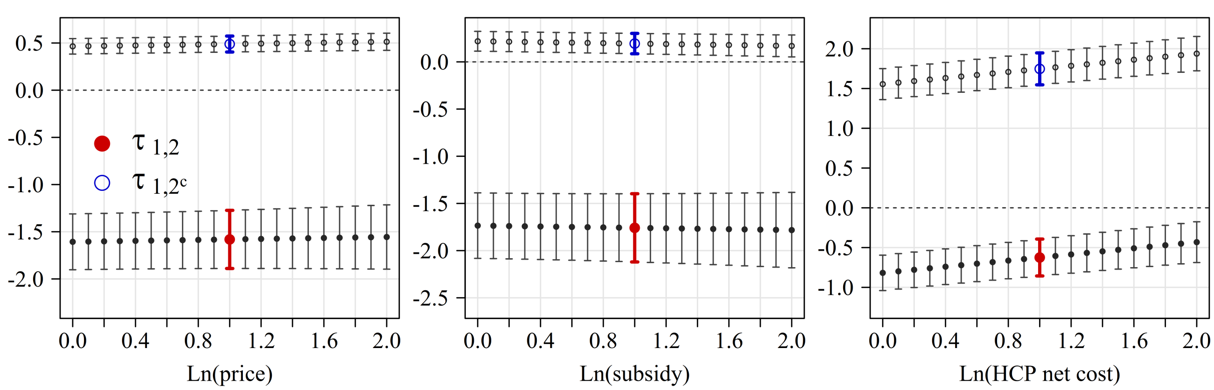

Figure 6: Testing the parallel-trends assumption.

Notes: Robustness of the baseline estimates to a measured linear violation of parallel trends (Aryal et al., 2025; Manski and Pepper, 2018), where g scales the assumed pre-trend differential ($g = 1$ is the standard parallel-trends benchmark). The red and blue point estimates plot the baseline ($g = 1$) effects of switching from $\mathcal{P}_1$ to $\mathcal{P}_2$ and from $\mathcal{P}_1$ to $\mathcal{P}_2^c$, respectively, with 95% confidence whiskers; gray extensions trace the estimates as g varies over $[0, 2]$.

Figure 6 reports the results, with one panel for each of the three outcomes (price, subsidy, HCP net cost). Within each panel, the tests for $\tau_{1,2}$ and $\tau_{1,2^c}$ are overlaid so the two effects can be compared at a glance. The central point in each pair is the baseline estimate at $g = 1$, and the gray extensions trace the robust estimate as g varies in $[0, 2]$. The estimates at $g = 1$ reproduce the baseline pattern. $\tau_{1,2} < 0$ and $\tau_{1,2} < \tau_{1,2^c}$. As g varies in $[0, 2]$, all the estimates retain their significance and economic magnitude. The $\tau_{1,2^c} - \tau_{1,2}$ gap is sizable and persistent across the full range of g, indirectly corroborating cross-subsidization.

5.6 Endogeneity in program selection

A natural concern is that the estimated treatment effects reflect selection rather than the causal impact of program design. HCPs choose which mechanism to join, and this choice may correlate with unobserved time-varying characteristics that also affect outcomes, violating the parallel-trends assumption that our ATT estimates require. We do not claim to eliminate endogeneity, but several features of the setting and the results collectively establish it as a second-order concern.

Since the three programs share identical eligibility criteria and administrative procedures, a rational HCP chooses among them by minimizing its net cost. An HCP therefore prefers $\mathcal{P}_2$ to $\mathcal{P}_1$ if its 35% cost share falls below the urban benchmark, i.e., if its 2013 cost ratio satisfies $p_{i,2013}/p_{u,i,2013} < 1/0.35 \approx 2.86$. We accordingly define the high-cost-ratio indicator $H_i = \mathbf{1}[p_{i,2013}/p_{u,i,2013} > 1/0.35]$, equal to one for HCPs whose 2013 cost ratio favors staying in $\mathcal{P}_1$. Because H_i is observable and time-invariant, HCP fixed effects absorb it, and the logit reported below confirms that observable selection is in fact concentrated in H_i. The residual identification requirement is therefore *parallel trends*. Do future switchers and future stayers have the same outcome trajectory absent the policy change?

Table 6 reports a logit of the 2014 switching decision on H_i and HCP-level covariates (ln(speed), ln(price), ln(requests)). H_i is negative and highly significant across four specifications with increasingly rich controls, and adding the additional covariates raises pseudo R^2 only from 0.031 to 0.037. Selection into switching is therefore concentrated in H_i, which the HCP fixed effects absorb. Whether selection on H_i creates non-parallel trends in u_{it} is addressed by the 2012 placebo in Section 5.4.

TABLE 6—LOGIT ESTIMATES OF THE SWITCHING DECISION FROM $\mathcal{P}_1$ TO $\mathcal{P}_2$ IN 2014.

	(1)	(2)	(3)	(4)
H_i	−0.966	−0.990	−1.094	−1.079
	(0.207)	(0.208)	(0.226)	(0.229)
ln(speed)		−0.180	−0.280	−0.283
		(0.086)	(0.121)	(0.122)
ln(price)			0.194	0.176
			(0.159)	(0.166)
ln(requests)				0.108
				(0.275)
Constant	−0.336	−0.091	−1.744	−1.595
	(0.169)	(0.204)	(1.373)	(1.427)
N	565	565	565	565
Pseudo R^2	0.031	0.037	0.037	0.036

Notes: Sample of the 565 HCPs on $\mathcal{P}_1$ in 2013 who continue into 2014, estimating the binary decision to switch to $\mathcal{P}_2$. All covariates are measured in 2013. $H_i = \mathbf{1}[p_i/p_{u,i} > 1/0.35]$ flags 2013 plans whose rural-to-urban price ratio exceeds the $\mathcal{P}_2$ break-even. Coefficients are log-odds. Pseudo R^2 is McFadden's adjusted statistic. Standard errors in parentheses.

The 2012 placebo in Section 5.4 directly tests parallel trends, and the placebo-to-headline ratios in Table 5 bound the residual selection bias. The DML estimates in Section 5.3 agree with the linear TWFE baseline, ruling out nonlinearity in the observables (though not selection on unobservables). The Aryal–Manski sensitivity in Section 5.5 shows the estimates survive a one-parameter family of departures from parallel trends. Appendix robustness checks confirm survival to influential observations (Appendix H), to a common-support restriction on the $\mathcal{P}_2^c$-relevant speed range (Appendix I), and to the Oster (2019) stability bounds (Appendix J). Endogeneity may shift precise magnitudes (and the placebo identifies two outcome-arm combinations for which this is more likely), but the direction and economic significance of the results survive every robustness exercise we performed.

5.7 Combined estimates over 2014–2021

The 2013–2014 design provides the cleanest identification in the panel. It is the only year-pair with a true policy change, and the parallel-trends assumption is directly testable via the 2012 placebo. The $\mathcal{P}_2$ arm rests on a comfortable 419 first-mover switchers. The $\mathcal{P}_2^c$ arm, resting on only 43, is too thin to anchor inferences on its own. To broaden the $\mathcal{P}_2^c$ evidence base, we estimate the same continuous-share TWFE canonical DiD on each year pair $(t-1, t)$ for $t \in \{2014, \ldots, 2021\}$.

Panel A of Figure 3 shows the gradual reallocation of RHC subsidy outlays across the three mechanisms from 2013 to 2021. $\mathcal{P}_1$'s share steadily shrinks while $\mathcal{P}_2$ and $\mathcal{P}_2^c$ expand. Program rules are stable from 2014 onward, so the variation that identifies later year-pair estimates is within-HCP movement between mechanisms. Treatment is non-absorbing, so the estimates average over both into-treatment and out-of-treatment movements.

For each $t \in \{2014, \ldots, 2021\}$ we restrict the sample to HCPs that appear in both $t-1$ and t, recover the six switching shares from the $\mathcal{C}_1 \cap \mathcal{C}_2$ identification scheme, and estimate the same continuous-share TWFE regression as in Section 5.3. The eight year-pair estimates of $\tau_{1,2}$, $\tau_{1,2^c}$, and their difference are reported in Table A9 and plotted in Figure 7 with 95% confidence ribbons. Year-pair estimates of $\tau_{1,2}$ are consistently negative, supporting the combined estimate as a representative summary of the program's recurring effects.

To summarise the year-pair evidence in a single number suitable for policy claims, we combine the eight year-pair estimates by inverse-variance weighting within each outcome. We call the resulting weighted average the *combined estimate* and report it in Table 7. Because some HCPs appear in multiple year pairs, the combined standard error understates the true sampling variance.

The combined estimate of $\tau_{1,2}$ on subsidy is -0.897, implying that switching from $\mathcal{P}_1$ to $\mathcal{P}_2$ reduces the subsidy bill by approximately 59%. The corresponding combined estimates on price and HCP net cost are -0.866 and -0.376 (58% and 31% reductions). For $\tau_{1,2^c}$ the picture is sharply different. Switching to $\mathcal{P}_2^c$ delivers only a 25% subsidy reduction (versus 59% for $\mathcal{P}_2$) and *increases* HCP net cost by 74% relative to $\mathcal{P}_1$ ($\hat{\tau}_{1,2^c}^{\text{net cost}} = +0.551$). The differential on net cost,

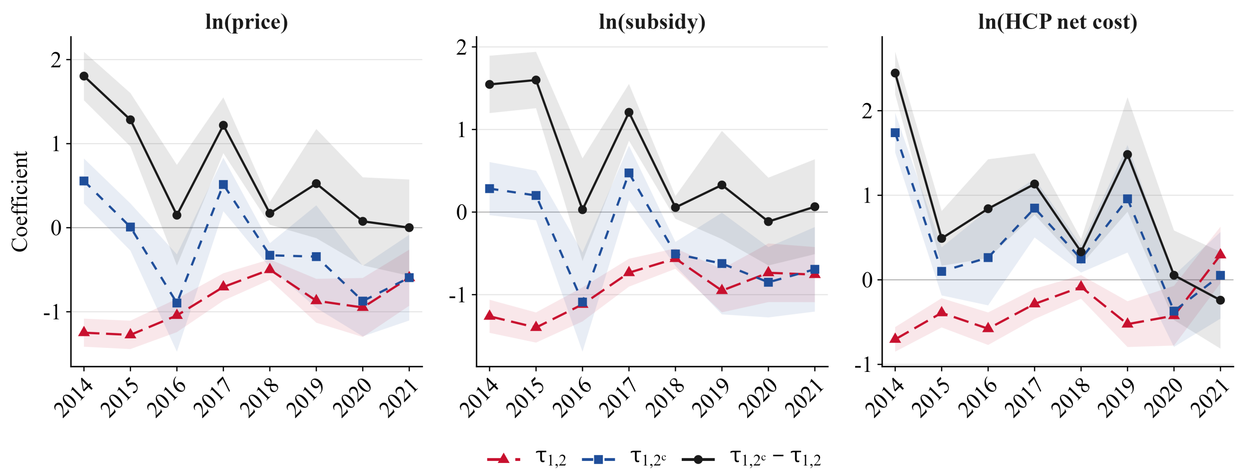

Figure 7: Year-pair continuous-share TWFE coefficients, 2014–2021.

Notes: Each panel plots three series with 95% confidence ribbons: $\tau_{1,2}$ ($\mathcal{P}_1 \to \mathcal{P}_2$, red), $\tau_{1,2^c}$ ($\mathcal{P}_1 \to \mathcal{P}_2^c$, blue), and the difference $\tau_{1,2^c} - \tau_{1,2}$ (dark gray). Estimates and standard errors are from Table A9.

$\hat{\tau}_{1,2^c} - \hat{\tau}_{1,2} = +0.816$, isolates the consortium-versus-individual gap.

TABLE 7—COMBINED TWFE CONTINUOUS-SHARE ESTIMATES, 2014–2021.

	ln(price)	ln(subsidy)	ln(HCP net cost)
$\tau_{1,2}$	−0.866	−0.897	−0.376
	(0.033)	(0.035)	(0.034)
$\tau_{1,2^c}$	−0.137	−0.290	0.551
	(0.051)	(0.053)	(0.053)
$\tau_{1,2^c} - \tau_{1,2}$	0.583	0.449	0.816
	(0.052)	(0.054)	(0.055)

Notes: Inverse-variance combined estimates of the year-pair TWFE continuous-share treatment effects reported in Table A9. For each outcome and each parameter ($\tau_{1,2}$, $\tau_{1,2^c}$, and the difference $\tau_{1,2^c} - \tau_{1,2}$), the combined estimate is $\hat{\tau}_{\text{comb}} = \sum_t \hat{\tau}_t / \hat{\sigma}_t^2 \big/ \sum_t 1/\hat{\sigma}_t^2$ with $\hat{\sigma}^2_{\text{comb}} = 1/\sum_t 1/\hat{\sigma}_t^2$, taken over the eight year pairs $t = 2014, \ldots, 2021$. The combined standard error treats the year-pair estimates as independent; because some HCPs appear in multiple year pairs, the reported standard error understates the true sampling variance, so the combined estimate should be interpreted as a summary reference for the year-pair trajectory rather than as a stand-alone inferential statistic. Standard errors in parentheses. Significance:

These estimates yield two policy claims. First, the 2021 RHC program disbursed about \$557 million (USAC, 2021), of which 43% went to $\mathcal{P}_1$, for a $\mathcal{P}_1$ subsidy base of roughly \$240 million. Applying $\hat{\tau}_{1,2}^{\text{subsidy}} = -0.897$ to this base implies that migrating all $\mathcal{P}_1$ HCPs to $\mathcal{P}_2$ would save about \$142 million annually, the cost-containment dividend of restoring demand elasticity predicted in Section 3. This estimate aligns with the \$143 million annual savings projected by Rabbani (2024a) from a two-part pricing model. Second, applying the differential $\hat{\tau}_{1,2^c}^{\text{subsidy}} - \hat{\tau}_{1,2}^{\text{subsidy}} = +0.449$ to the 2021 $\mathcal{P}_2^c$ subsidy total of about \$194 million implies that \$70 million of $\mathcal{P}_2^c$ outlays (36% of the arm, or 13% of the program budget) leaks from eligible to ineligible consortium members through consortium-internal price allocation. Discontinuing $\mathcal{P}_2^c$ eliminates this leakage at no cost to eligible HCPs, since $\mathcal{P}_2$ delivers a strictly cheaper net-cost outcome on average.

5.8 An empirical test of cross-subsidization

Section 5.7 found that $\mathcal{P}_2^c$ prices exceed $\mathcal{P}_2$ prices, suggestive of cross-subsidization. To test the mechanism, we exploit a sharper theoretical prediction. Cross-subsidization requires a mix of eligible and ineligible members, so the eligible-member price distortion must peak at an interior mix. This implies an inverted U-shape between eligible members' price and the ineligible fraction.

To test this hypothesis, we limit the data to consortium applications ($\mathcal{P}_2^c$). So far, the empirical analysis has used only eligible HCPs. Here we identify ineligible members and measure their share of each consortium. We identify ineligible members as HCPs who (1) are part of a consortium application that has been approved and some members of the consortium have been reimbursed, (2) have internet speed on file, and (3) have received no subsidy.

We use two measures of each consortium's ineligible share. The main measure is speed-based. It is the fraction of the consortium's total speed belonging to ineligible members, with speed serving as a proxy for internet consumed. The second measure is the headcount share, for example 0.40 when two of five consortium members are ineligible. We combine consortium-application data across all available years to obtain a sufficient sample, with summary statistics in Table A8. On average, consortium members are 21.7% ineligible by headcount and 19.6% by speed.

We apply Frisch–Waugh–Lovell (FWL) residualization to isolate variation in the ineligible fraction that is orthogonal to high-dimensional controls. We regress ln(price) and the ineligible fractions on the mean number of bidders in the auction, the natural logarithm of the average internet speed requested per eligible member, the natural logarithm of the sum of all speeds requested by the consortium (a measure of consortium size), and year and consortium FE. We then apply local linear regression to the residuals using a nonparametric LOWESS (locally weighted scatterplot smoothing) estimator and express the result on the original-variable scale by adding back the means. The procedure flexibly estimates the conditional relationship between the ineligible fraction and ln(price) for eligible members without imposing a parametric functional form.[11]

Figure 8 reveals the inverted-U shape predicted by Proposition **P.3**, with a peak near $R \approx 0.2$–0.3. The interior peak rules out a selection-on-size alternative, which would predict a monotone relationship. This adds a third strand of evidence for cross-subsidization, joining the within-HCP difference $\tau_{1,2^c} - \tau_{1,2}$ in Section 5.3 and the institutional record in Section 2.

[11] Implemented in R via *ggplot* with *geom_smooth* method set to *loess*.

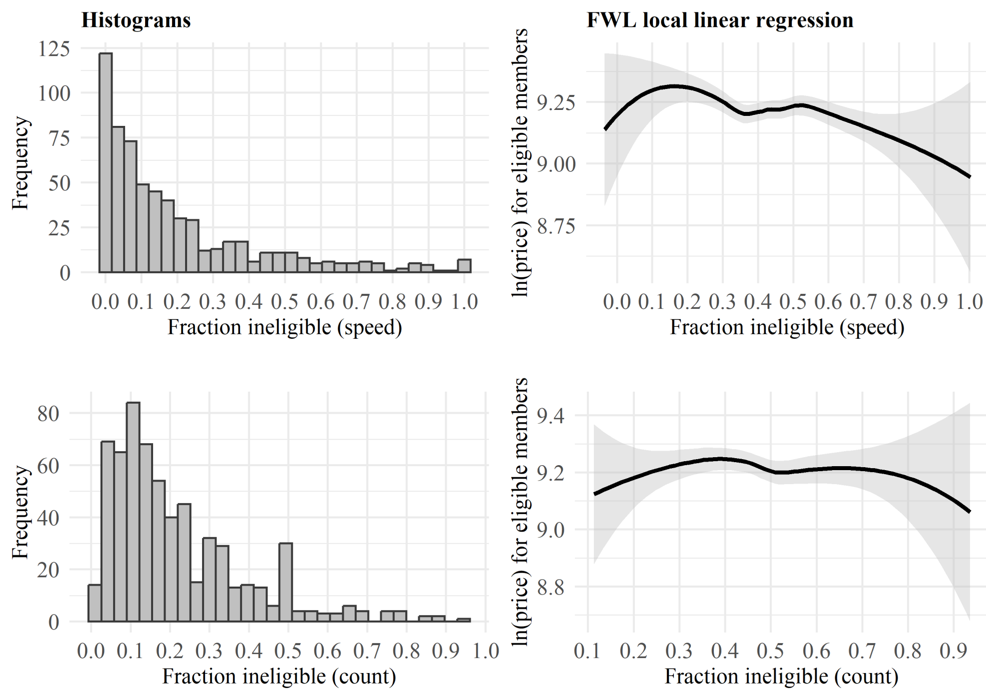

Figure 8: FWL local linear regression results.

Notes: Sample pools consortium-application data across all available years. Left panels: histograms of the consortium ineligible fraction by speed share (top) and headcount share (bottom). Right panels: Frisch–Waugh–Lovell residualized LOESS fits of ln(price) for eligible members on the ineligible fraction, with 95% confidence bands, partialling out consortium and year fixed effects, ln(total consortium speed), ln(mean speed), and the mean number of bidders.

6 Discussion

Our results are consistent with the cost-effectiveness hypothesis. Combining year-pair estimates over 2014–2021, switching from $\mathcal{P}_1$ to $\mathcal{P}_2$ delivered a 58% reduction in internet price, 59% in subsidy, and 31% in HCP net cost. Migrating all $\mathcal{P}_1$ users to $\mathcal{P}_2$ would extend these subsidy and net-cost reductions program-wide.

We also found evidence consistent with the cross-subsidization hypothesis. Switching from $\mathcal{P}_1$ to $\mathcal{P}_2^c$ reduces subsidies by only 25%, far less than the 59% reduction delivered by $\mathcal{P}_1 \to \mathcal{P}_2$. Applied to the 2021 $\mathcal{P}_2^c$ base of $194 million, this gap implies roughly $70 million per year of excess outlays (Section 5.7). The contrast is sharper for HCP net cost. $\mathcal{P}_1 \to \mathcal{P}_2$ delivered a 31% reduction, while $\mathcal{P}_1 \to \mathcal{P}_2^c$ delivered a 74% increase. That HCPs voluntarily switched to a materially costlier program is hard to explain without latent incentives such as cross-subsidization.

The original intention of the USF was to ensure that rural HCPs do not pay more than their urban counterparts. Even setting aside cross-subsidization, $\mathcal{P}_2^c$'s extension of subsidies to urban HCPs appears unjustified. Majority-rural status is determined by a head count, so a large urban hospital can secure 65% subsidies by forming a consortium with two small rural clinics. By extending subsidies to urban HCPs, $\mathcal{P}_2^c$ has expanded the recipient base and strained the program budget.

We therefore recommend discontinuing $\mathcal{P}_1$ and $\mathcal{P}_2^c$ and migrating all HCPs to $\mathcal{P}_2$. The three mechanisms share identical eligibility criteria, so the population with subsidized connectivity is invariant to the choice of mechanism. Migrating $\mathcal{P}_1$ to $\mathcal{P}_2$ delivers the same services at lower government cost while HCPs gain on average ($\hat{\tau}_{1,2}^{\text{net cost}} = -0.376$), since the 35% copayment's negotiation incentive more than offsets the higher marginal cost share. ISPs lose the rents that $\mathcal{P}_1$'s design enabled. Discontinuing $\mathcal{P}_2^c$ eliminates a transfer with no policy rationale: the \$70 million per year routed to ineligible consortium members falls outside the USF statute's stated targets. Applying a marginal cost of public funds of at least 1.20 (Ballard, Shoven, and Whalley, 1985), the combined savings translate into \$254 million in annual social welfare gains via reduced deadweight loss of taxation.

Cross-subsidization adds to the literature on subsidies that target one group and exclude another. Rotemberg (2019) documents the crowding-out of ineligible rivals. We find the mirror image, a positive spillover from eligible to ineligible entities within consortia. Policymakers designing such programs must consider not only the direct effect on intended recipients but also indirect effects on third parties.

Subsidy-induced cost manipulation has been documented in other USF programs. In the High Cost program, the loop support pays a percentage of reported investment costs, with higher percentages for higher reported costs. Berg, Jiang, and Lin (2011a,b) document that firms inflate reported investment costs to qualify for higher subsidies, with inflation rising in the anticipated subsidy increment. Both findings reflect the same design vulnerability of subsidy formulas tied to self-reported metrics.

The RHC has several facets worth investigating that lie beyond the scope of this study. First, the program's \$400 million budget cap went unmet for two decades and was common knowledge. An unmet cap can signal to recipients the availability of unallocated funds, incentivizing aggressive subsidy pursuit and inviting wasteful spending and cost overstatement (Berg, Jiang, and Lin, 2011a).

A second avenue, building directly on our study, is the role of market concentration in cross-subsidization. Cross-subsidization requires sustained cooperation between an ISP and a consortium, which is harder to sustain when more parties are involved (Bourreau, Sun, and Verboven, 2021; Dal Bó and Fréchette, 2018; Sugaya and Wolitzky, 2025). Consistent with this, Clemens and Gottlieb (2017) document that Medicare subsidies inflate prices more in concentrated markets. Higher concentration among ISPs, HCPs, or both, may therefore intensify the cost burden of cross-subsidization in the RHC.

A third avenue concerns the RHC's impact on connectivity and rural healthcare, which existing studies have not addressed. Other USF programs offer cautionary parallels. The E-Rate program increased school connectivity by 68% but did not improve student outcomes (Goolsbee and Guryan, 2006; Hazlett, Schwall, and Wallsten, 2019). The High Cost program delivers less benefit than its subsidy expense for many North Carolina recipients (Boik, 2017). Most of the Lifeline program's budget went to existing users rather than new connectivity (Ackerberg et al., 2014; Lyons, 2023; Wallsten, 2016), and three-quarters of eligible low-income households did not take up the subsidy while many ineligibles received it (Ward and Woroch, 2010).

A subsidy program is unjustified unless it achieves its intended impact. The RHC aims to promote HCP internet connectivity, but subsidized HCPs earn combined annual revenues exceeding \$1 trillion.[12] Internet subsidies account for less than 0.1% of HCPs' operating scale, making it implausible that the average HCP would forgo connectivity if the subsidy were withdrawn. If the program fails to boost connectivity, it reduces to a wealth transfer from taxpayers to HCPs and the ISPs that serve them.

7 Conclusions

This paper provides a comprehensive economic evaluation of the Federal Communications Commission's Rural Health Care Program (RHC), analyzing three subsidy mechanisms: price-cap ($\mathcal{P}_1$), individual ad valorem subsidy ($\mathcal{P}_2$), and consortium ad valorem subsidy ($\mathcal{P}_2^c$). Using administrative data and causal econometric models, we find that transitioning from $\mathcal{P}_1$ to $\mathcal{P}_2$ reduces government subsidy spending by 59% and lowers what HCPs pay for internet, while preserving the program's effectiveness.

In contrast, the consortium option enables an unintended cross-subsidization scheme in which ineligible members use eligible members as a conduit to receive subsidies. Eligible members bear inflated internet prices, mirrored by deflated prices for ineligible members, driving up the consortium's subsidy outlays. Our estimates are consistent with the presence of cross-subsidization. The existing enforcement regime, limited to occasional audits and modest fines, lacks the incentives to contain it.

The RHC's costs have been rising for three decades. Concerns about inefficiencies and loopholes are documented (Section A.1), but little has been done to fix them. Instead, policymakers have focused on raising more funds and taxing new entities to keep the program afloat (FCC, 2012; Figliola, 2025; Kelly et al., 2025). Our study provides a rigorous economic diagnosis that has been missing, along with actionable recommendations. Open questions for future research include the role of market concentration in cross-subsidization, the program's effects on connectivity and health outcomes, and sophisticated forms of cross-subsidization involving ISPs.

Data and code availability. The data used in this study are administrative records from the Federal Communications Commission's Universal Service Administrative Company, publicly available at opendata.usac.org. The code that produces every table and figure in the paper is available from the authors upon request.

[12] In 2023, US hospitals generated \$1.5 trillion. 67.4% of these hospitals are non-profit or public entities.

Online appendix. A separate online appendix (`00_online_appendix.pdf`) contains institutional details, the theoretical model appendix (proofs and extensions), data cleaning, machine-learning implementation, additional robustness tables, placebo tests, Oster bounds, influence diagnostics, and common-support analysis.

References

Ackerberg, Daniel A., David R. DeRemer, Michael H. Riordan, Gregory L. Rosston, and Bradley S. Wimmer. 2014. "Estimating the impact of low-income universal service programs." *International Journal of Industrial Organization* 37:84–98. URL https://doi.org/10.1016/j.ijindorg.2014.07.009.

Aryal, Gaurab, Charles Murry, Pallavi Pal, and Arnab Palit. 2025. "The Benefits from Bundling Demand in K-12 Broadband Procurement." Working Paper 33498, National Bureau of Economic Research. URL https://doi.org/10.3386/w33498.

Ballard, Charles L., John B. Shoven, and John Whalley. 1985. "General Equilibrium Computations of the Marginal Welfare Costs of Taxes in the United States." *American Economic Review* 75 (1):128–138.

Berg, Sanford V, Liangliang Jiang, and Chen Lin. 2011a. "Incentives for cost shifting and misreporting: US rural universal service subsidies, 1991–2002." *Information Economics and Policy* 23 (3-4):287–295. URL https://doi.org/10.1016/j.infoecopol.2011.09.001.

Berg, Sanford V., Liangliang Jiang, and Chen Lin. 2011b. "Universal service subsidies and cost overstatement: Evidence from the U.S. telecommunications sector." *Telecommunications Policy* 35 (7):583–591. URL https://doi.org/10.1016/j.telpol.2011.04.016.

Boik, Andre. 2017. "The economics of universal service: An analysis of entry subsidies for high speed broadband." *Information Economics and Policy* 40:13–20. URL https://doi.org/10.1016/j.infoecopol.2017.04.004.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *Review of Economic Studies* 91 (6):3253–3285. URL https://doi.org/10.1093/restud/rdae007.

Bourreau, Marc, Yutec Sun, and Frank Verboven. 2021. "Market entry, fighting brands, and tacit collusion: Evidence from the French mobile telecommunications market." *American Economic Review* 111 (11):3459–3499. URL https://doi.org/10.1257/aer.20190540.

Caliendo, Marco and Steffen Künn. 2011. "Start-up subsidies for the unemployed: Long-term evidence and effect heterogeneity." *Journal of Public Economics* 95 (3-4):311–331. URL https://doi.org/10.1016/j.jpubeco.2010.11.003.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2024. "Difference-in-Differences with a Continuous Treatment." *Journal of Econometrics* URL https://doi.org/10.3386/w32117. Forthcoming; NBER Working Paper 32117.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. "Double/debiased machine learning for treatment and structural parameters." *The Econometrics Journal* 21 (1):C1–C68. URL https://doi.org/10.1111/ectj.12097.

Cicala, Steve, Ethan MJ Lieber, and Victoria Marone. 2019. "Regulating markups in US health insurance." *American Economic Journal: Applied Economics* 11 (4):71–104. URL https://doi.org/10.1257/app.20180011.

Clemens, Jeffrey and Joshua D Gottlieb. 2017. "In the shadow of a giant: Medicare's influence on private physician payments." *Journal of Political Economy* 125 (1):1–39. URL https://doi.org/10.1086/689772.

Conkling, Thomas S. 2018. "Crowd-Out or Affordability? The Lifeline Expansion's Effect on Wireless Service Spending." *Journal of Policy Analysis and Management* 37 (2):357–383. URL https://doi.org/10.1002/pam.22053.

Dal Bó, Pedro and Guillaume R Fréchette. 2018. "On the determinants of cooperation in infinitely repeated games: A survey." *Journal of Economic Literature* 56 (1):60–114. URL https://doi.org/10.1257/jel.20160980.

de Chaisemartin, Clément and Xavier d'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110 (9):2964–2996. URL https://doi.org/10.1257/aer.20181169.

Decarolis, Francesco, Maria Polyakova, and Stephen P. Ryan. 2020. "Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D." *Journal of Political Economy* 128 (5):1712–1752. URL https://doi.org/10.1086/705550.

Eriksson, Ross C, David L Kaserman, and John W Mayo. 1998. "Targeted and untargeted subsidy schemes: Evidence from postdivestiture efforts to promote universal telephone service." *The Journal of Law and Economics* 41 (2):477–502. URL https://doi.org/10.1086/467398.

Fairlie, Robert W and Jonathan Robinson. 2013. "Experimental evidence on the effects of home computers on academic achievement among schoolchildren." *American Economic Journal: Applied Economics* 5 (3):211–240. URL https://doi.org/10.1257/app.5.3.211.

FCC. 2012. "Universal Service Contribution Methodology; A National Broadband Plan for Our Future: Further

Notice of Proposed Rulemaking." Tech. Rep. FCC 12-46, FCC, Washington, D.C. URL https://docs.fcc.gov/public/attachments/FCC-12-46A1.pdf. WC Docket No. 06-122; GN Docket No. 09-51.

Figliola, Patricia Moloney. 2025. "The Universal Service Fund and Related FCC Broadband Programs: Overview and Considerations for Congress." CRS Report R47621, Congressional Research Service. URL https://www.congress.gov/crs_external_products/R/PDF/R47621/R47621.9.pdf.

Fremeth, Adam, Brian Kelleher Richter, and Brandon Schaufele. 2018. "Spillovers from regulating corporate campaign contributions." *Journal of Regulatory Economics* 54:244–265. URL https://doi.org/10.1007/s11149-018-9369-7.

Geruso, Michael and Timothy Layton. 2020. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." *Journal of Political Economy* 128 (3):984–1026. URL https://doi.org/10.1086/704756.

Goolsbee, Austan and Jonathan Guryan. 2006. "The impact of Internet subsidies in public schools." *The Review of Economics and Statistics* 88 (2):336–347. URL https://doi.org/10.1162/rest.88.2.336.

Greenstein, Shane. 2020. "The basic economics of internet infrastructure." *Journal of Economic Perspectives* 34 (2):192–214. URL https://doi.org/10.1257/jep.34.2.192.

Gruber, Jonathan and Ebonya Washington. 2005. "Subsidies to employee health insurance premiums and the health insurance market." *Journal of Health Economics* 24 (2):253–276. URL https://doi.org/10.1016/j.jhealeco.2004.08.004.

Hanna, Mark, Scott J. Savage, and Bradley Wimmer. 2025. "The Impacts of Subsidy Design on Rural Internet Access and Quality." Tech. Rep. SSRN 5087366, SSRN. URL https://doi.org/10.2139/ssrn.5087366.

Hazlett, Thomas W, Ben Schwall, and Scott Wallsten. 2019. "The educational impact of broadband subsidies for schools under E-rate." *Economics of Innovation and New Technology* 28 (5):483–497. URL https://doi.org/10.1080/10438599.2018.1527554.

Kelly, Mark, Markwayne Mullin, Mike Crapo, and Kevin Cramer. 2025. "Kelly, Mullin, Crapo, Cramer Introduce Bill to Expand High-Speed Internet Across the U.S." Press release, US Senate. URL https://www.kelly.senate.gov/newsroom/press-releases/kelly-mullin-crapo-cramer-introduce-bill-to-expand-high-speed-internet-across-the-u-s/. "Lowering Broadband Costs for Consumers Act".

Lyons, Daniel. 2023. "Assessing Broadband Affordability Initiatives." Tech. rep., American Enterprise Institute. URL https://lira.bc.edu/works/publication-article/f9jk4-94267.

Manski, Charles F and John V Pepper. 2018. "How do right-to-carry laws affect crime rates? Coping with ambiguity using bounded-variation assumptions." *Review of Economics and Statistics* 100 (2):232–244. URL https://doi.org/10.1162/rest_a_00689.

Mendez, Samara, Gabor Molnar, and Scott J Savage. 2021. "The Impacts of the Lifeline Subsidy on High-Speed Internet Access." *The Journal of Law and Economics* 64 (4):745–782. URL https://doi.org/10.1086/714504.

Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics* 37 (2):187–204. URL https://doi.org/10.1080/07350015.2016.1227711.

Rabbani, Maysam. 2024a. "Dollars and megabits: a comparative analysis of Telecom and Healthcare Connect Fund." *Information Economics and Policy* 67:101082. URL https://doi.org/10.1016/j.infoecopol.2024.101082.

———. 2024b. "Internet price, speed, and disparity: The case of rural healthcare providers in the United States." *Telecommunications Policy* 48 (2):102674. URL https://doi.org/10.1016/j.telpol.2023.102674.

Rotemberg, Martin. 2019. "Equilibrium effects of firm subsidies." *American Economic Review* 109 (10):3475–3513. URL https://doi.org/10.1257/aer.20171840.

Sugaya, Takuo and Alexander Wolitzky. 2025. "Repeated Games with Many Players." Working paper.

USAC. 2021. "Universal Service Fund Annual Report 2021." Tech. rep., USAC, Washington, DC. URL https://www.usac.org/about/reports-orders/annual-report/.

———. 2024. "Universal Service Fund Annual Report 2024." Tech. rep., USAC, Washington, DC. URL https://www.usac.org/about/reports-orders/annual-report/.

Vernon, Richard. 1979. "Unintended consequences." *Political Theory* 7 (1):57–73. URL https://doi.org/10.1177/009059177900700104.

Wallsten, Scott. 2011. "The Universal Service Fund: What Do High-Cost Subsidies Subsidize?" Tech. rep., Technology Policy Institute. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1927933.

———. 2016. "Learning from the FCC's Lifeline Broadband Pilot Projects." Tech. Rep. TPRC 44, Technology Policy Institute. URL https://techpolicyinstitute.org/wp-content/uploads/2016/03/Wallsten_Learning-from-the-FCCs-Lifeline-Broadband-Pilot-Projects.pdf.

Ward, Michael R and Glenn A Woroch. 2010. "The effect of prices on fixed and mobile telephone penetration: Using price subsidies as natural experiments." *Information Economics and Policy* 22 (1):18–32. URL https://doi.org/10.1016/j.infoecopol.2009.12.001.