Estimating the Price Elasticity of Demand of the Premium on New Enrollment in the ACA Marketplaces

Justin Li

Abstract:

The U.S. Affordable Care Act's health insurance Marketplaces rely on open-enrollment periods (OEP) to attract new consumers, yet the extent to which premiums affect enrollment is unclear. We aim to find the price elasticity of demand for these plans. We have two hypotheses: 1) demand is elastic, so the percentage change in new consumers exceeds the percentage change in average premiums, or 2) demand is inelastic, so the opposite holds. We compare a baseline model to progressively more rigorous models, our final model being a fixed effects model that controls for time trends that affect all states equally and state-specific characteristics. The baseline estimate suggests a moderately elastic response of –0.8, meaning a 1 percent premium increase predicts a 0.8 percent drop in new enrollment. However, introducing state fixed effects reduces the elasticity to –0.4 percent, and adding time fixed effects further narrows down the elasticity to –0.2 percent, which we found is not statistically significant. Much of the observed premium-enrollment relationship in pooled data is explained by state-specific characteristics, which could include a given states' regulations, insurer competition, or outreach efforts, rather than changes in the premium. We conclude that while a change in the premium can influence enrollment, its impact is statistically insignificant once a greater context is accounted for. Policymakers should therefore aim to combine premium subsidies or rate regulation with state-specific strategies. These strategies can include targeted efforts to subsidy design, competition, or outreach to more effectively boost ACA Marketplaces' enrollment.

Introduction:

Background

Health insurance is the primary means by which households get timely access to care and protect themselves against catastrophic medical costs. Studies have found that insured adults have higher rates of preventive-service use, fewer avoidable hospitalizations, and better long-term health outcomes than their uninsured peers (Riedel, 2009; Hermer, 2005). Yet coverage is incomplete in the United States. Early-estimate data from the National Health Interview Survey show that in the first half of 2024, roughly 26 million people, or 7.9 percent of the population, lacked insurance at the time of interview (Briones & Cohen, n.d.).

In the U.S., a household (the “buyers”) pays a premium to an insurer. When care is needed, the insurer then pays most of the bill that hospitals, physicians, and drug makers (the “sellers”) charge, while patients pay a smaller share for the service. Because prices are negotiated in advance, the arrangement pools risk and smooths spending over time. But it also creates agency problems; for example, patients may use more care once they are insured (Morrisey, 2008).

The key pieces of a health-insurance contract are the premium, deductible, coinsurance and copay, and the out-of-pocket maximum. The premium is the fixed monthly payment for insurance, and the deductible is the amount the enrollee must spend (usually per year) on covered care before the plan begins sharing the costs. Coinsurance is the percentage that the enrollee pays after meeting the deductible, while copay is the flat fee that the enrollee pays after meeting the deductible. Finally, the out-of-pocket maximum is the ceiling on what the enrollee pays in a year. Once it is reached, the insurer pays all of the covered costs (Morrisey, 2008).

2023 Current Population Survey data show that 92.0 percent of the nation’s 331.7 million residents (~305 million) had health insurance at some point during 2023. Private coverage, led by employer-sponsored plans, remained the dominant form of insurance, enrolling 65.4 percent of the population, while public programs such as Medicare and Medicaid covered 36.3 percent. None of these percentages changed meaningfully from 2022 (U.S. Census Bureau, 2024). The ten largest carriers, led by UnitedHealth Group, Elevance (Blue Cross Blue Shield), CVS (Aetna), and Cigna, account for more than half of national enrollment.

Households are willing to pay a premium that exceeds their expected medical costs in exchange for certainty. However, potential issues arise. The first is moral hazard: lower service prices because of insurance lead patients to consume more care than they would if they faced the full price. Adverse selection is another issue: people with higher expected costs are more likely to seek more coverage, threatening the pool’s average cost (Morrisey, 2008).

Scope and Objectives

Health-insurance markets are characterized by wide variability in plan design and price, but relatively little is known about how those attributes shape the quantity of insurance demanded in the United States. This paper addresses that gap by seeking to answer how the

design of health insurance plans affects plan enrollment in the United States. Several factors of insurance influence coverage, and we focus specifically on the premium. We look at what happens to the number of users of insurance with the increase in premiums, holding all else fixed. Specifically, our objective is to determine the percentage change of new consumers within a state for each percentage increase from the previous year's average premium. This will help us estimate the price elasticity of demand.

The price elasticity of demand tells us how sensitive buyers are to price. The premium is the "price," and the number of new enrollees is the "quantity" demanded. A value of –0.5, for example, would mean a 1 percent premium increase is linked to a 0.5 percent drop in new enrollment within that state (Pendzialek, Simic, & Stock, 2016).

We have two hypotheses. The first is that the price elasticity of demand is elastic: The percent change in new consumers of Marketplace Open Enrollment Period (OEP) plans is greater than the percent change in the year's average premium, after controlling for confounding factors like COVID-19 impacts, general time trends, state specific time trends, and state-specific non time varying factors.

The other hypothesis is that the price elasticity of demand is inelastic: The percent change in new consumers of Marketplace OEP plans is less than the percent change in the average premium, after controlling for confounding factors like COVID-19 impacts, general time trends, state-specific time trends, and state-specific non time varying factors.

Literature Review

There is a gap in the recent literature on the impact of insurance plan design on the quantity of insurance users. Chernew, Cutler, and Keenan measure how rising health-insurance premiums during the 1990s led to less coverage among non-elderly adults in 64 U.S. metropolitan areas. Using Population Survey data linked to local premiums, they used probit and instrumental-variables models that look at variation in premium growth, controlling for labor market and Medicaid changes. They found that more than half of the decline in coverage rates experienced over the 1990s is attributable to the increase in health insurance premiums (2.0 percentage points of the 3.1 percentage point decline) (Chernew, Cutler, & Keenan, 2005).

Baicker and Chandra look at how higher premiums affect jobs. They use the natural experiment of the early-2000s spike in malpractice costs that pushed premiums up. Applying a difference-in-differences method, they found that a 10 percent jump in premiums cuts employment by 1.6 percent, working hours by 1 percent, and moves about 2 percent of workers into part-time roles. Wages for employees who stay covered fell by roughly 2.3 percent, showing that employers pass on only part of the higher costs (Baicker & Chandra, 2005).

Phelps and Newhouse show how the details of an insurance plan change how much health care people use. They model two prices using pre-RAND data: the coinsurance rate and the time cost of getting care. Higher coinsurance cuts use, but the size of that cut shrinks when a visit already takes a lot of the patient's time. Their study shows that both time and money

matter, and that cost-sharing rules shape behaviour once someone is insured (Phelps & Newhouse, 1974).

In 2014, a systematic review of empirical studies on price elasticity of demand for health insurance in managed-competition settings finds that U.S. studies report elasticities ranging from –0.2 to –1.0 for optional primary coverage (which includes the ACA Marketplaces). This range shows that there is substantial variation across contexts and suggests that while some populations are relatively inelastic, others respond strongly to price changes (Pendzialek, Simic, & Stock, 2016).

Methods:

Data

The data we used were the Open Enrollment Period (OEP) State-Level Public Use File (PUF) from the Marketplace OEP PUFs from CMS.gov. We used eight PUFs for each year from 2017–2024, giving us n=408 total observations between eight years across 50 states plus the District of Columbia.

Each observation had relevant variables such as the number of consumers and new consumers, average premium, important demographics, and the number of consumers in each plan level. Demographics included age categories, race, and state. The plan levels, called metal levels, were Catastrophic, Bronze, Silver, Gold, and Platinum in increasing cost, coverage, and actuarial value (the percent of the cost of a plan that the insurer expects to cover). Not every state Marketplace offered Catastrophic or Platinum plans. Additionally, some observations lacked average premiums.

We controlled for several variables, namely each of the metal levels and demographics. Using a fixed effects model accounts for possible omitted variable bias when the variables remain constant over time. We assumed that demographic variables, namely age, sex, and race, and the distribution between the five metal levels would vary over time, hence the inclusion of the variables as control variables.

Our analyses looked at how the Average Premium (defined as the average monthly premium per person before the application of APTC for all consumers with a Qualified Health Plan) affects the number of New Consumers (defined as the count of unique consumers who selected a Qualified Health Plan that did not have Qualified Health Plan coverage in the previous year, including those who were enrolled in a Qualified Health Plan of an HC.gov state or a different SBM (State Based Marketplace) state in the previous year).

Possible confounding variables include time- and state-based differences. COVID-19 was a potent example of what we looked to control for. Other possible confounders were state policy changes and Medicaid expansion. Subsidies, specifically the Advance Premium Tax Credit–a federal tax credit to help eligible consumers pay their monthly health insurance premiums–could be possible moderators in the relationship between premiums and new consumers (Centers for Medicare & Medicaid Services, 2025).

Understanding the fixed effects model

The data we use is panel data, meaning each state has data across a specific time period (2017–2024). Some differences are permanent to each state, and others come from certain events that affect everyone at the same time, in the same way. Using a fixed-effects model cleans those patterns out, helping us isolate the relationship between premiums and enrollment.

An intuitive way to understand the fixed effects is that it gives every state its own starting line. The "state fixed effect" absorbs factors that remain constant for that state. Namely, it accounts for size, culture, laws, etc.; California's naturally huge baseline enrollment and Alaska's naturally small enrollment are taken into account. We need to control for these differences. States with large enrollment on average also have large premiums on average; states with low new enrollment on average have small premiums on average. If this is true, a pooled regression would be biased upwards. The model also considers the shared effects that each state experiences each year. These "time fixed effects" factors in nationwide events that affect all states the same way, like nationwide legislation changes, COVID-19, etc.

The model aims to capture the remaining effect of the premium on enrollment. Fixed effects strip away the constant traits of each state and time-specific factors. We need to control for these differences. As a hypothetical example: states with large enrollment on average also have large premiums on average, and states with low new enrollment on average have small premiums on average. If this is true, a pooled regression would be biased upwards. Enrollment growth could also be following a trend over time. This would change the amount of influence that the premium would have on enrollment. Controlling for those effects gives us insight on how premium and enrollment are related causally.

Model

We ran four models: a base model (no fixed effects), a state fixed effects model, a time fixed effects model, and a two-way fixed effects (time and state) model.

The **full formal model** (Two-way Fixed Effects) is described below:

$$\log(Y_{it}) = \beta_0 + \beta_1(\log(\text{premium}_{it})) + \beta_2(\text{state}_i) + \beta_3(\text{year}_t) + \beta_4(X_{it}) + \varepsilon_{it}$$

Where i = state, and t = time. X_{it} is a vector of time varying control variables on the state level. These include metal plan, age, and race distribution over time within each state. The Y_{it} is the predicted total number of new enrollees within a state over time. The β_1 is the predicted percentage change of the number of new enrollees when the premium changes by 1 percent within a state over time.

Three other models are described below:

Base model:

$$\log(Y_{it}) = \beta_0 + \beta_1(\log(premium_{it})) + \varepsilon_{it}$$

State Fixed Effects:

$$\log(Y_{it}) = \beta_0 + \beta_1(\log(premium_i)) + \beta_2(state_{it}) + \beta_3(X_{it}) + \varepsilon_{it}$$

Time Fixed Effects:

$$\log(Y_{it}) = \beta_0 + \beta_1(\log(premium_t)) + \beta_2(year_{it}) + \beta_3(X_{it}) + \varepsilon_{it}$$

The Two-way Fixed Effects model was chosen based on the Bayesian Information Criterion (Table 1).

All analyses were performed in R statistical software.

Results:

Descriptive Results

The following are descriptive figures that look at trends in premium and enrollment across the 50 states plus the District of Columbia from 2017–2024.

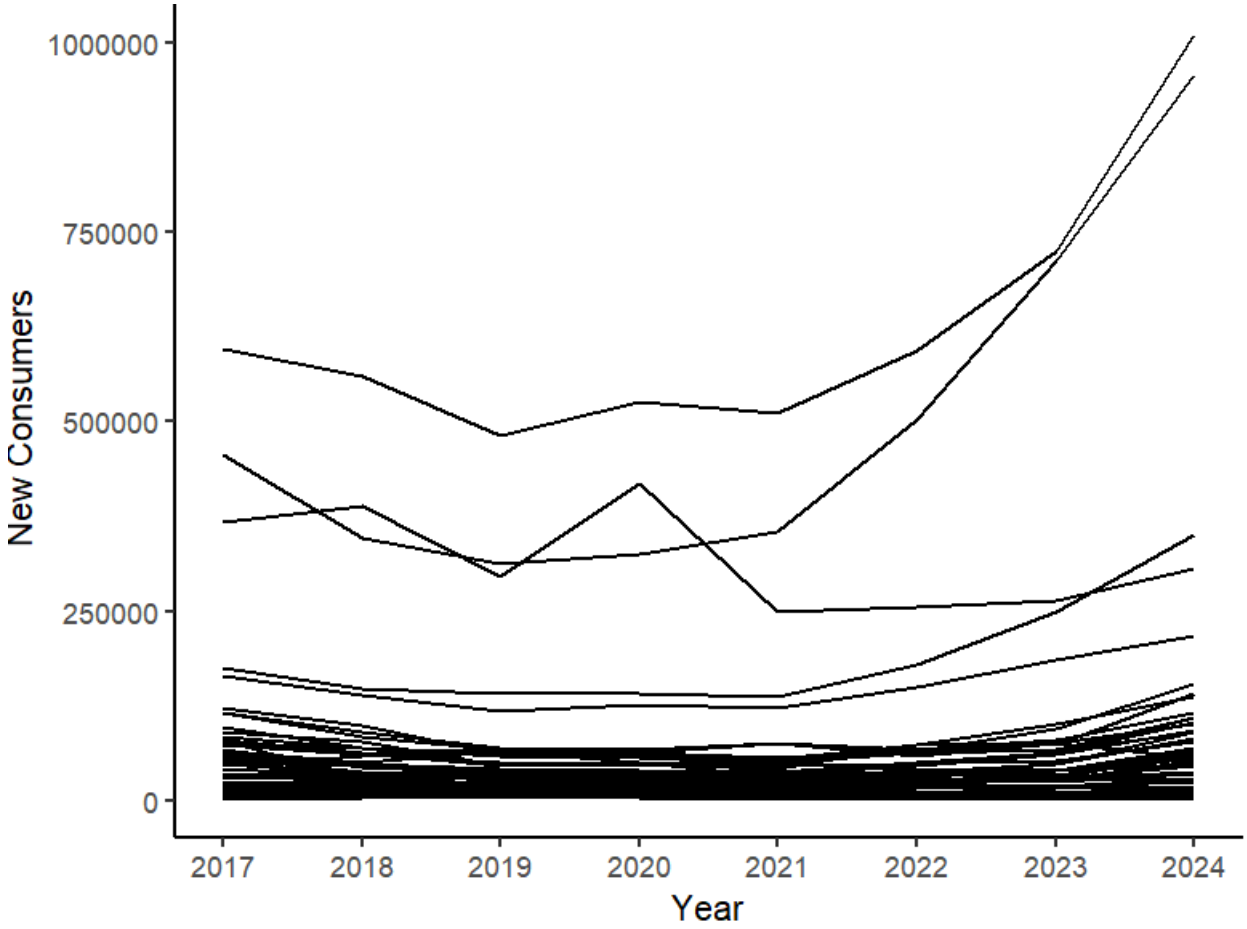

Figure 1. Profile plot of New Consumers among N=408 total observations across 50 states plus the District of Columbia from 2017–2024. New Consumers are defined as the count of unique consumers who selected a Qualified Health Plan that did not have Qualified Health Plan coverage in the previous year. New Consumers include those who were enrolled in a Qualified Health Plan of an HC.gov state or a different SBM (State Based Marketplace) state in the previous year.

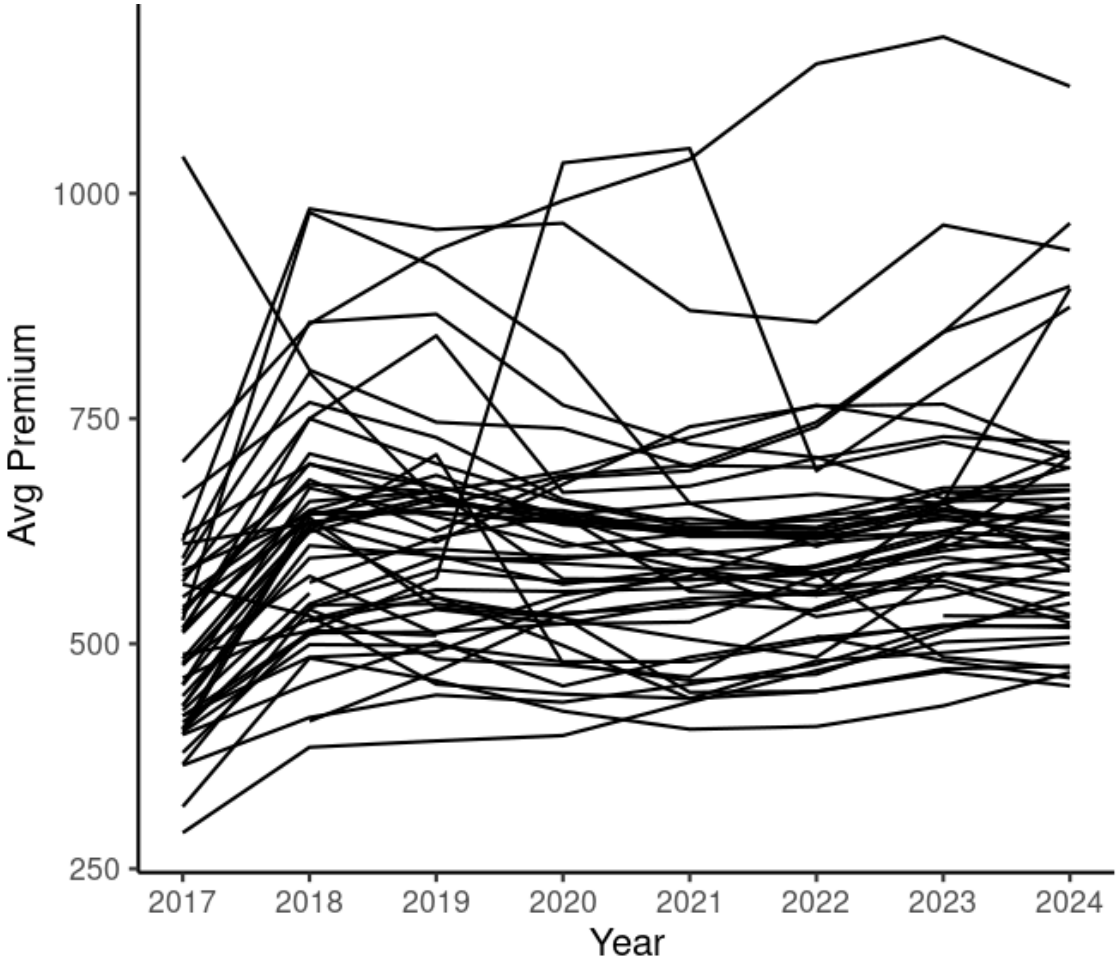

Figure 2. Profile plot of Average Premium (average monthly premium per person before the application of APTC for all consumers with a Qualified Health Plan) [cite CMS.gov Public Use Files Definitions] among N=398 observations across 50 states plus the District of Columbia from 2017–2024. There are N=408 total observations; there are 10 observations without Average Premiums.

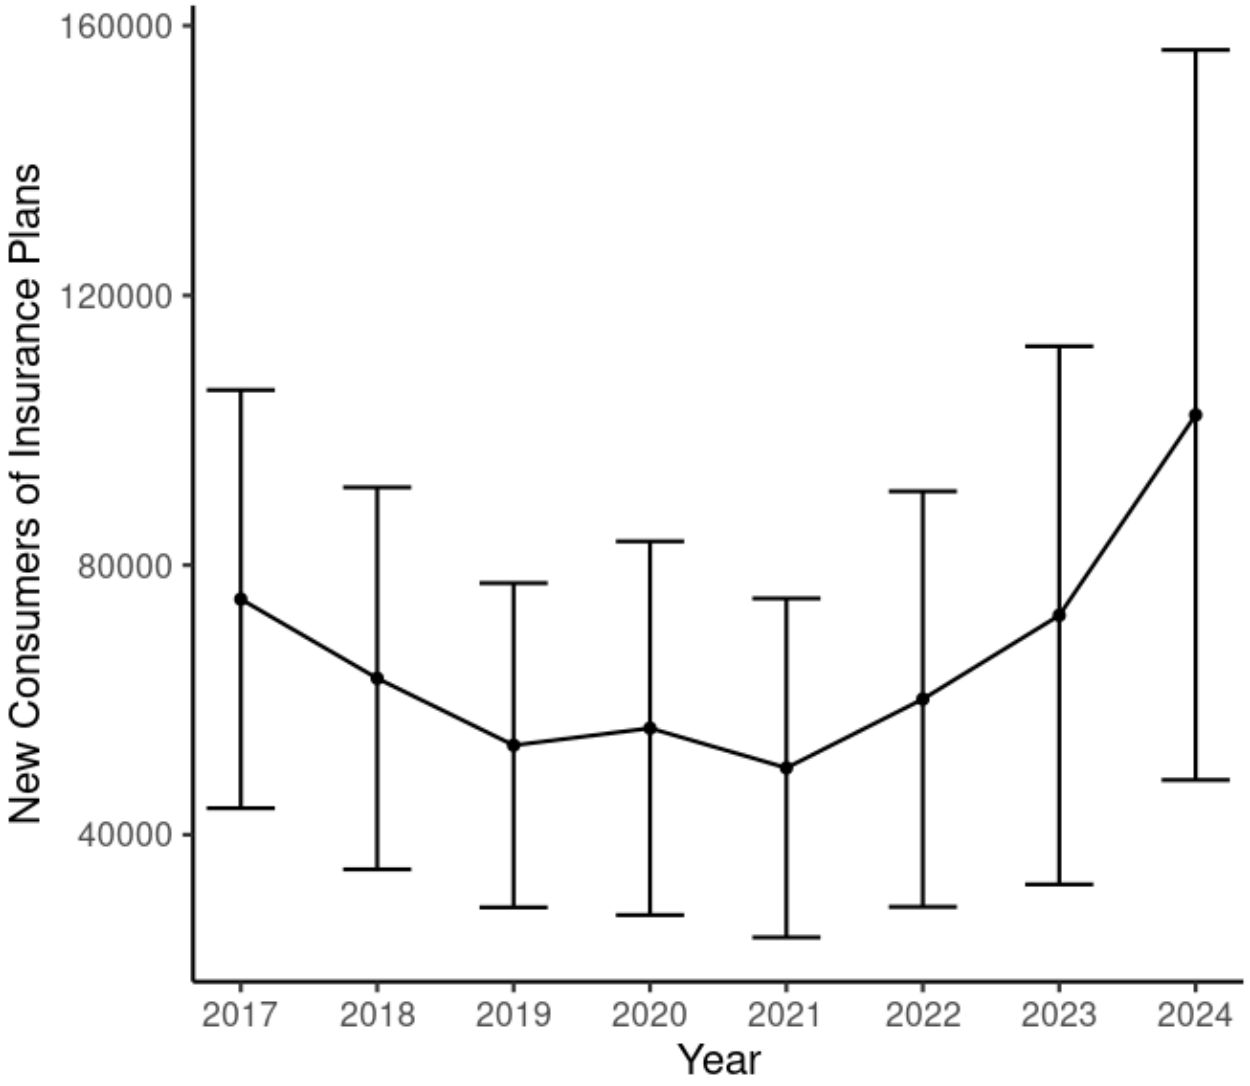

Figure 3. Empirical summary plot of New Consumers vs Time among 50 states plus the District of Columbia. The population average at each time labeled with black filled circle. Population trend shown by black line. Error bars mean ±2SE.

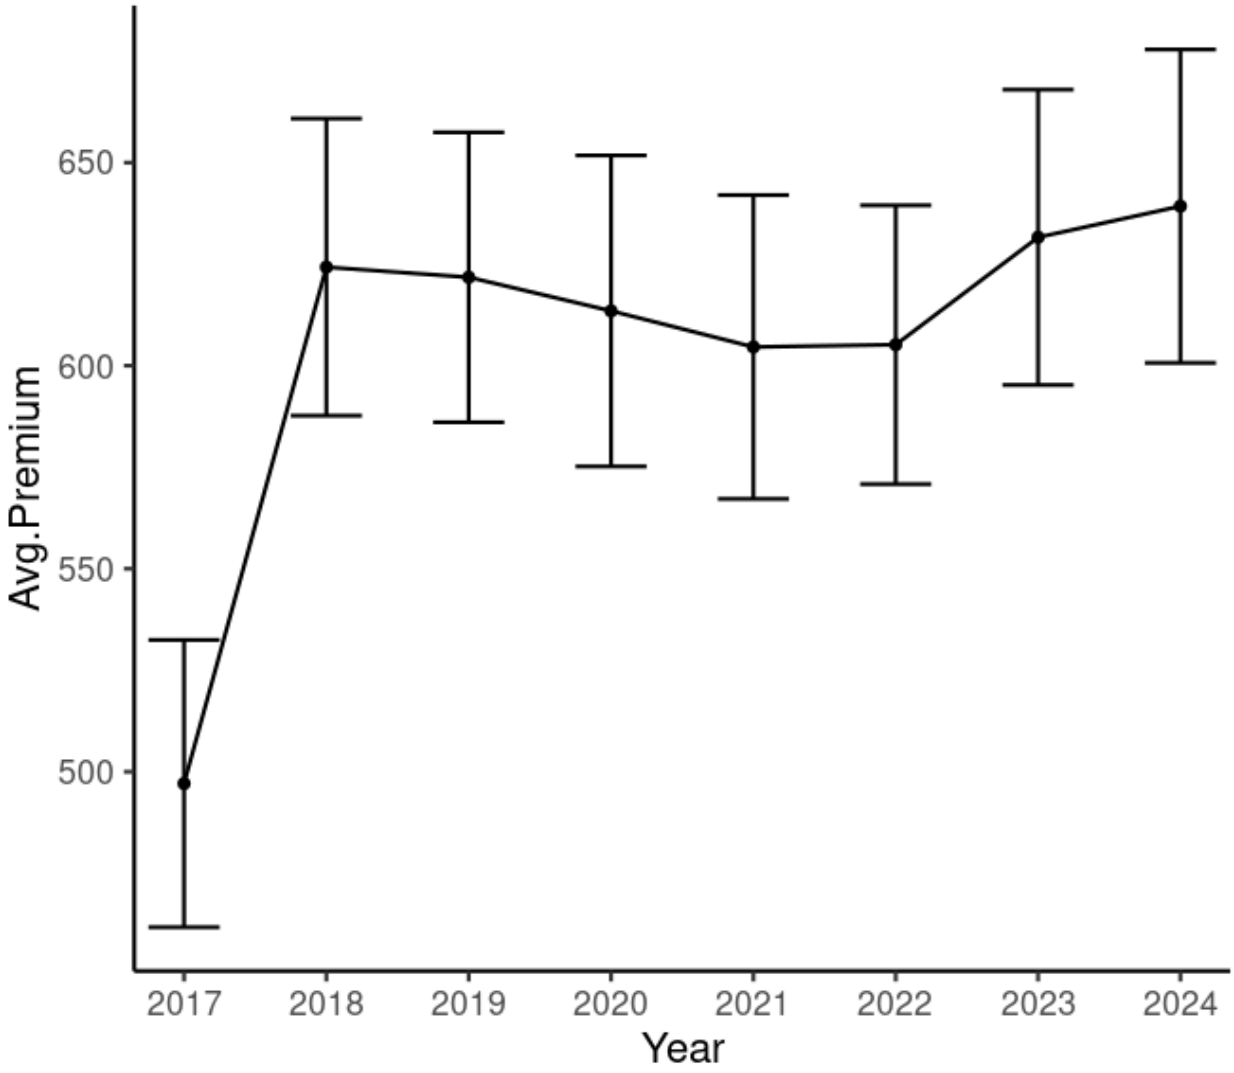

Figure 4. Empirical summary plot of Average Premium vs Time among 50 states plus the District of Columbia. The population average at each time labeled with black filled circle. Population trend shown by black line. Error bars mean ±2SE.

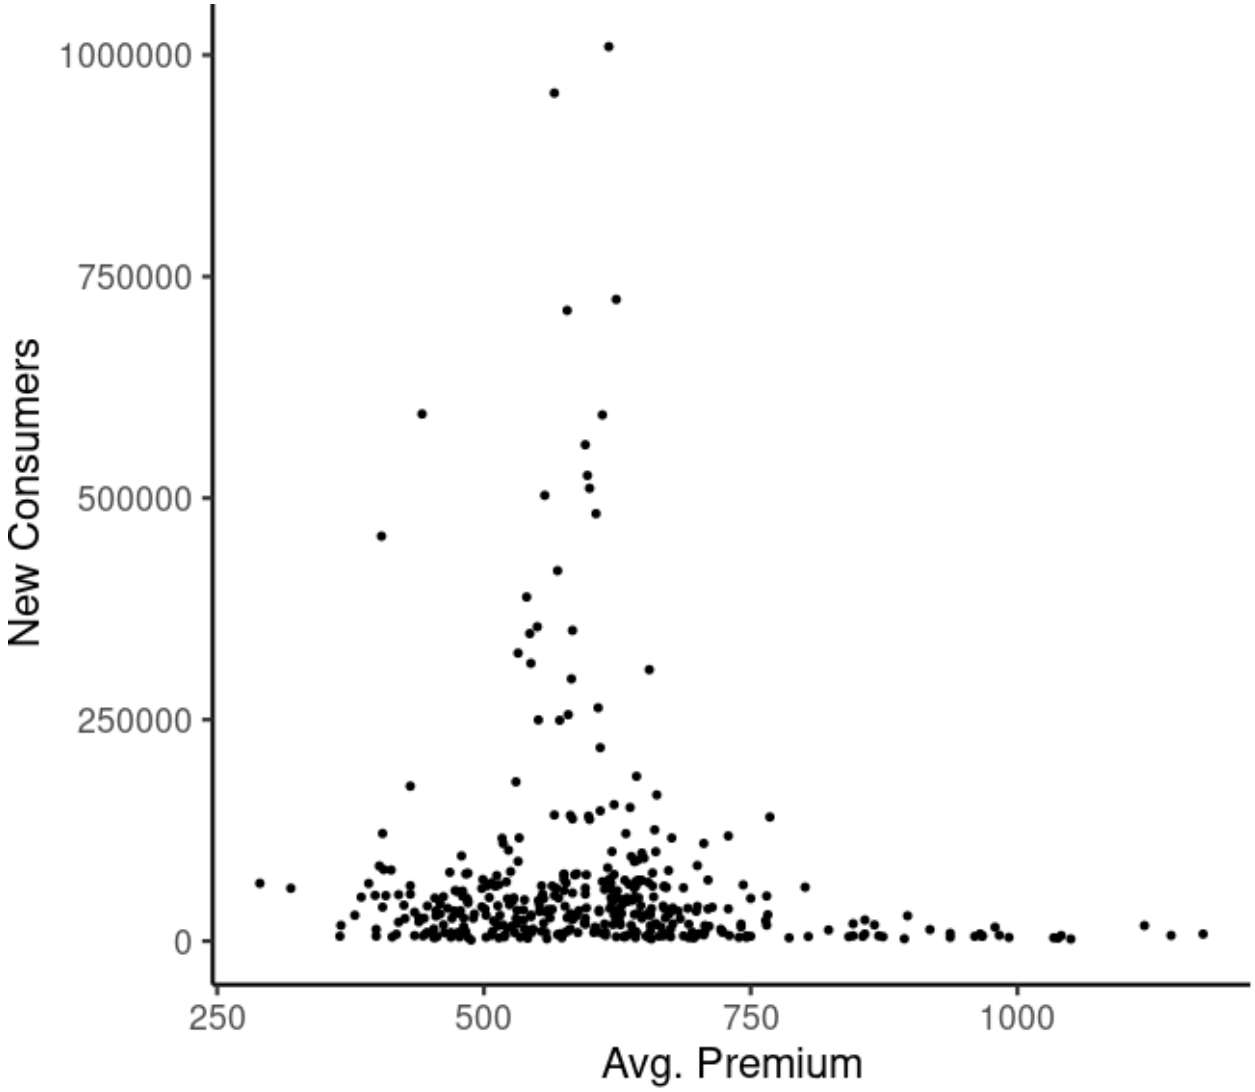

Figure 5. Plot of Average Premium and New Consumers among N=398 observations across 50 states plus the District of Columbia from 2017–2024. There are N=408 total observations; there are 10 observations without Average Premiums.

Model Results

Table 1. Comparing model fit using Bayesian Information Criterion (BIC)

Model	BIC
Base Model	629
State Fixed Effects	354
Time Fixed Effects	625
Two-way Fixed Effects	194

Notes: The BIC evaluates the overall model fit, considering all variables, including the fixed effects. It determines which model best explains the change in the new customers, balancing fit against complexity. Lower BIC means a better overall model (ScienceDirect Topics, n.d.).

Table 2. The Impact of Premiums on Health Insurance Marketplace Enrollment, 2017–2024

Model	Elasticity Estimate	Standard Error
Base Model	0.8 ***	0.2
State Fixed Effects	0.4 **	0.1
Time Fixed Effects	0.7 ***	0.2
Two-way Fixed Effects	0.2	0.1

*Notes: *** p<0.001, ** p<0.01, * p<0.05. Analyses were performed in R software. All models control for metal level enrollment (Bronze, Silver, Gold), demographics (race, gender, age), Marketplace type (SMB and HC.gov), and total Insurance beneficiaries. Final data contained N=309 observations (original N=408) across 50 states plus the District of Columbia from 2017-2024. Some panels were unbalanced. State fixed effects control for the unique characteristics of states that remain constant over time. Time fixed effects control for general time trends that affect all states in a similar way. Two-way fixed effects includes both state and time fixed effects. The final model's (two-way fixed effects) elasticity estimate reflects the percentage change in new consumer enrollment for a Health Insurance Marketplace plan selected during the Open Enrollment Period with a 1% change in average premium over time. This estimate is compared to what it would be given the overall nationwide time trend (time fixed effects) and state-specific averages (state fixed effects), while holding constant metal level enrollment (Bronze, Silver, Gold), demographics (race, gender, age), and Insurance beneficiaries.*

Our updated model in Table 2 indicates that the price elasticity of insurance enrollment was a lot smaller than we initially estimated without controlling for time and state. Whereas the baseline model had an elasticity of –0.8, the final model with state fixed effects and time fixed effects estimated a much smaller elasticity of –0.2, which is statistically insignificant.

Much of the relationship between price and enrollment that appears originally seems to be explained by state-level characteristics. Examples of these characteristics could be

demographics or policy climate specific to certain states. The original model predicted a –0.8 elasticity; with state fixed effects –0.4 elasticity.

Controlling for those state-specific effects gave a less significant elasticity; controlling for the effects of general time trends decreased it even more. Our two-way fixed effects model found a –0.2 price elasticity of demand for the premium, a statistically insignificant effect on enrollment.

Discussion

Our findings show that once we account for each state's characteristics–like its regulations and market structure–the strong raw link between premiums and enrollment largely fades. California, for instance, could consistently enroll far more people than Alaska, regardless of short-run premium movements, because its basic conditions make it inherently have higher coverage. After we sweep those conditions out with state fixed effects, a premium increase is associated with only a modest fall in new enrollment, indicating that price changes alone have less impact than the pooled comparison would suggest.

State factors are important because they shape the entire Marketplace. Subsidies, outreach, and insurer competition have the potential to cushion price shocks, while weaker policy support, poorer populations, or insurer concentration can magnify them. Because of this, the elasticity of demand is unlikely to be uniform across states. In high-subsidy states, consumers may be less affected by premium increases; in concentrated markets with few plan options, they may react sharply because they cannot switch to a cheaper alternative. It would be interesting to look at premiums after subsidies, or controlling for finer state traits, for future research.

Conclusion

Once we account for unobserved state characteristics and general time trends, the price elasticity of demand for ACA Marketplace OEP plans is statistically zero. While our original pooled model suggested a –0.8 percent elasticity, adding state fixed effects halved that estimate (–0.4 percent), and incorporating time fixed effects reduced it further to –0.2 percent, which we found was statistically insignificant. These results imply that cross-state differences in demographics, subsidy generosity, insurer competition, and outreach efforts drive most of the observed enrollment variation, rather than within-state premium fluctuations.

The practical lesson is that cutting premiums through subsidies or other methods could raise enrollment, but probably less than policymakers hope. States with price-sensitive consumers would gain the most from premium assistance, while states whose residents appear price-inelastic might see bigger returns from non-price tools such as outreach campaigns or broader provider networks. Any nationwide premium policy should therefore be paired with a strategy tailored to each state's specific characteristics.

In short, premiums matter, but the state context matters more. Treating premium regulations or cuts as a silver bullet is likely to not be effective unless complemented by policies

that take into account other factors, like subsidy design, competition, outreach, etc., that influence whether households take up coverage or not.

Limitations

Our study uses the listed premiums for each plan, even though many consumers are eligible for the Advance Premium Tax Credit that lowers what they actually pay. Because we do not take that into account, namely, use the average premium after APTC, we may overstate their out-of-pocket premium and understate how strongly subsidized buyers react to price changes.

The average premium is also affected by enrollment. When new people join a market, they change the average itself. This endogeneity can decrease the estimated relationship between premiums and coverage, even though controlling for metal-level reduces the problem. The data are at the state-year level, so we cannot track how people move between individual plans with different deductibles, networks, or actuarial values. Plan-level data would give a clearer view of these choices. This would be beneficial for follow-up research.

Finally, we had to drop 99 state-year observations (about 20 percent of the sample) because of missing premium or enrollment figures, or missing control variable data, such as metal level data. If these gaps are random, they mainly reduce precision; however, if there is clustering in states that experienced policy shifts, this could result in biased results.

References

1. Baicker, K., & Chandra, A. (2005). Effects of rising malpractice costs on health insurance and employment. *Journal of Health Economics, 24*(3), 673–682.

2. Briones, E. M., & Cohen, R. A. (n.d.). *Health insurance coverage: Early release of estimates from the National Health Interview Survey, January–June 2024*. Centers for Disease Control and Prevention. www.cdc.gov/nchs/data/nhis/earlyrelease/insur202412.pdf

3. Centers for Medicare & Medicaid Services. (2025). *Qualified Health Plan Coverage and Premium Public Use File (PUF)*. Retrieved August 3, 2025, from cms.gov/cciio/resources/data-resources/marketplace-puf

4. Chernew, M. E., Cutler, D. M., & Keenan, J. R. (2005). Impact of rising health-insurance premiums on coverage rates in U.S. metropolitan areas. *Health Services Research, 40*(4), 1021–1039. https://doi.org/10.1111/j.1475-6773.2005.00409.x

5. Hermer, L. D. (2005). Private health insurance in the United States: A proposal for a more functional system. *Houston Journal of Health Law & Policy, 6*, 1–84.

6. Morrisey, M. A. (2008). *Health insurance* (2nd ed.). Health Administration Press.

7. Phelps, C. E., & Newhouse, J. P. (1974). Coinsurance, time costs, and the demand for medical services. *Journal of Health Economics, 13*(2), 151–169. https://doi.org/10.2307/1923971

8. Pendzialek, J. B., Simic, D., & Stock, S. (2016). Differences in price elasticities of demand for health insurance: A systematic review. *European Journal of Health Economics, 17*(1), 5–21. https://doi.org/10.1007/s10198-014-0650-0

9. Riedel, L. M. (2009). Health insurance in the United States. *AANA Journal, 77*(6), 439–444.

10. ScienceDirect Topics. (n.d.). *Bayesian information criterion (BIC)*. Retrieved August 3, 2025, from www.sciencedirect.com/topics/medicine-and-dentistry/bayesian-information-criterion#:~:text=Bayesian%20information%20criterion%20(BIC)%20(

11. U.S. Census Bureau. (2024). *Health insurance coverage in the United States: 2023* (Current Population Reports P60-284). U.S. Government Publishing Office.