

The Subsidy Cliff and the Risk Pool: An Actuarial Analysis of Adverse Selection and Premium Stability in the ACA Individual Market after the Expiration of the Enhanced Premium Tax Credits

Oluwatosin Dorcas Babalola¹, Oluwakemi Oyinlade², Daniel Ekow Arthur³

¹Maurice R. Greenberg School of Risk Science, J. Mack Robinson College of Business, Georgia State University, Atlanta, GA, USA.

²Maurice R. Greenberg School of Risk Science, J. Mack Robinson College of Business, Georgia State University, Atlanta, GA, USA.

³Maurice R. Greenberg School of Risk Science, J. Mack Robinson College of Business, Georgia State University, Atlanta, GA, USA.

ABSTRACT

Background. The enhanced Affordable Care Act premium tax credits expired on 31 December 2025, restoring the original contribution schedule and the cliff at 400% of the federal poverty level.

Methods. Net premiums were computed for every HealthCare.gov rating area from CMS public use files. A Tweedie cost model was estimated on 68,119 MEPS person-years and a selection-and-repricing model was calibrated to 2026 plan selections. Income-band responses were tested against 2024 to 2025 as a placebo, in state-by-band difference-in-differences regressions with the unsubsidized premium burden as a dose, and in a county panel. Loss ratios for 2021 to 2024 and the cost per person-year of four restoration schedules were computed.

Results. The median benchmark premium rose 24.3%; crossing 400% of poverty costs a 40-year-old \$1,487 and a 60-year-old \$10,163. In 30 HealthCare.gov states, selections above 400% fell 34.6% against 3.5% below, after rising 9.0% in the placebo year. A single price response calibrated to the total predicts a 5.1% loss above the threshold; fitting the segments separately requires responses differing by a factor of 11. Selections at 300 to 400% of poverty exceed their counterfactual by 122,752, 71% of the loss immediately above. The exposed segment's loss rose 3.25 points per point of unsubsidized burden; the placebo year shows no gradient. Full restoration costs \$30,786 per person-year of coverage; extending the 9.96% cap above 400% costs \$9,636.

Conclusions. The subsidized market absorbed the reversion; the unsubsidized segment contracted sharply, partly through income reporting at the notch. No single elasticity describes both.

Keywords: Affordable Care Act; premium tax credits; adverse selection; subsidy cliff; bunching; difference-in-differences; medical loss ratio

This is an Open Access article that uses a fund-ingmodel which does not charge readers or their institutions for access and distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>) and the Budapest Open Access Initiative (<http://www.budapestopenaccessinitiative.org/read>), which permit unrestricted use, distribution, and reproduction in any medium, provided the original work is properly credited.

INTRODUCTION

The American Rescue Plan Act of 2021 lowered the share of income that subsidized households pay toward a benchmark marketplace plan and removed the ceiling at 400% of the federal poverty level above which the original Affordable Care Act pays no premium tax credit. The Inflation Reduction Act

extended both provisions through 2025. They expired on 31 December 2025, and the 2026 plan year was priced and enrolled under the original schedule.

Two literatures predict what should follow. Selection theory since Akerlof (1970) and Rothschild and Stiglitz (1976) holds that when a pooled price rises, the enrollees with the least

to gain from coverage leave first, raising the average cost of those who remain. Empirical work has found the mechanism in employer plans and in state markets (Cutler and Reber, 1998; Hackmann, Kolstad and Kowalski, 2015). The literature on notches holds that a schedule that withdraws a benefit entirely at a threshold induces bunching on the favorable side of it (Kleven and Waseem, 2013). The 2026 schedule contains both a price rise and a notch.

This paper asks five questions in the order the mechanism runs: how net premiums changed by income and age (RQ1); what the enrollment response implies for the pool's expected cost (RQ2); whether the subsidized and unsubsidized segments behaved as one market or two (RQ3); what insurers' loss ratios were going into the change and what the rebate floor protects (RQ4); and what restoring credits, in full or in part, would cost per person-year of coverage (RQ5).

The main finding concerns the structure of the response. A standard selection model calibrated to the aggregate decline in HealthCare.gov enrollment gives a small price response and a first-year morbidity rise of about 1%. The same model, asked to reproduce the two segments of the market separately, cannot do so with one parameter: the response above 400% of poverty is 11 times the response below it. Three pieces of evidence support reading this as a response to the notch rather than to price alone. Enrollment immediately below the threshold grew while every other subsidized band shrank, and did not do so in the prior year. Across states, the exposed segment shrank most where full-price coverage takes the largest share of income, and no such gradient existed in the prior year. The same pattern holds in a county panel with county-by-period fixed effects. The policy costing carries the structure through: a dollar of restored credit buys several times more coverage above the threshold than below it, and removing the cliff alone costs a fraction of full restoration. The paper quantifies these trade-offs and takes no position on them.

INSTITUTIONAL BACKGROUND

The credit

The premium tax credit equals the premium of the second-lowest-cost silver plan in the enrollee's rating area at the enrollee's age, less an expected contribution equal to household income times an applicable percentage that rises with income as a share of the poverty line. Because the credit moves with the benchmark, a subsidized enrollee's payment for the benchmark plan does not change when the gross premium rises; Jaffe and Shepard (2020) examine the consequences of this linkage for pricing. An enrollee without a credit pays every dollar of every increase. Premiums vary with age within a 3:1 band, with rating area and with tobacco use, and with nothing else.

The change

The 2025 schedule required nothing below 150% of poverty, rising to 8.5% of income at 400% and capped at 8.5% with no upper limit. The 2026 schedule is the original law as indexed

by Rev. Proc. 2025-25 (Internal Revenue Service, 2025): 2.10% of income below 133% of poverty, rising through 4.19% at 150%, 6.60% at 200% and 8.44% at 250%, to 9.96% from 300% to 400%, and no credit above 400%. Coverage in 2025 used the 2024 poverty guidelines (\$15,060 for one person in the contiguous states) and coverage in 2026 the 2025 guidelines (\$15,650); Alaska and Hawaii have their own. All calculations are for single adults.

Legislative status

The Senate voted on competing extension proposals on 11 December 2025 and neither reached 60 votes, so the enhanced credits expired on schedule and 2026 coverage was priced and sold without them. A later restoration would not reprice 2026 but would change what the counterfactuals in Section 6.6 represent.

RELATED LITERATURE

Einav, Finkelstein and Cullen (2010) and Einav and Finkelstein (2011) recast adverse selection as a demand curve meeting an average-cost curve that slopes down because the marginal buyer is cheaper than the average one; the repricing loop used here is that framework in actuarial form. Handel (2013) shows that reduced inertia can worsen selection, relevant to a year in which many enrollees faced an active decision. Geruso and Layton (2017) review selection and its remedies, including the risk adjustment model of Kautter et al. (2014). Finkelstein, Hendren and Shepard (2019) find that low-income adults' willingness to pay for coverage falls well below its cost and that enrollment responds strongly to the subsidized price; Saltzman (2019) estimates exchange demand in California and Washington. Published price responses differ in units and populations, so this paper treats its response grid as a sensitivity range and replaces it with values calibrated to 2026. Saez (2010) develops estimation from excess mass at kinks, and Kleven and Waseem (2013) show that a notch creates a dominated region above the threshold and predicts bunching below it. Marketplace eligibility is set on projected income at application and reconciled at tax filing, so the response available to a household near 400% of poverty includes a reporting decision. The cost model uses the Tweedie compound Poisson-gamma distribution, standard for claims with a point mass at zero and a long right tail (Smyth and Jørgensen, 2002).

DATA

Premiums

The CMS Exchange Rate and Plan Attributes public use files for plan years 2025 and 2026 list every plan's premium by age and rating area for states on HealthCare.gov: 31 states in 2025 and 30 in 2026, after Illinois moved to its own exchange. Child-only and SHOP plans are excluded. Issuers file a sentinel rate above \$9,000 a month for age and plan combinations they do not sell; these are removed before plans are ordered. The

benchmark is the second-lowest silver premium in each rating area at each age. Comparisons across years use the 349 rating areas with a benchmark in both.

Enrollment

The CMS Open Enrollment Period state-level public use files for 2024, 2025 and 2026 report plan selections by state, income band and age band, and the county-level files report the same with one band above 400% of poverty. Plan selections are not effectuated enrollment; CMS has not published 2026 effectuated enrollment as of the date above. Three features of the files affect every comparison. Each state file carries three national total rows. The 100 to 138% band is a subgroup of the 100 to 150% band and is excluded from sums, after which the income bands reconcile to the published totals. Platform totals compare different states across years because Georgia left HealthCare.gov for 2025 and Illinois for 2026, so every platform comparison here assigns each state its 2026 platform in all years. In addition, the 2026 file notes that Connecticut, Minnesota and New York reported enrollees with known incomes above 400% of poverty as “other income”, and the District of Columbia suppressed most income cells; all four report zero or near-zero selections above 400% in 2026 against 97,000 in 2025. Band-level analyses exclude them. The primary sample is the 30 HealthCare.gov states, which the premium data describe; a 47-state sample is reported for the national costing.

Cost

The Medical Expenditure Panel Survey full-year consolidated files for 2019 to 2024 provide annual total expenditure, coverage by month, age, sex, income relative to poverty, region and nine chronic-condition indicators. Expenditure is inflated to 2024 dollars and person weights are divided by the number of pooled years. MEPS identifies only 1,817 under-65 non-group person-years across six years, too few to price on, so the cost structure is estimated on all 68,119 under-65 privately insured person-years and the non-group indicator is tested as a covariate.

Insurer experience

The CMS Medical Loss Ratio public use files for reporting years 2021 to 2024 report earned premium, incurred claims, life-years and the statutory rebate calculation by issuer and state. The files carry resubmissions as separate rows; each year is deduplicated to the latest submission per issuer and state, and grand-total rows and territories are dropped, after which 2024 individual-market life-years total 22.43 million.

Section 1332 waivers

States operating a state-based reinsurance program under a Section 1332 waiver in 2025 and 2026 were taken from the CMS State Innovation Waivers page and its December 2022 data brief: 15 states, of which seven use HealthCare.gov (Alaska, Delaware, Montana, New Hampshire, North Dakota, Oregon and Wisconsin).

METHODS

Net premiums and the cliff (RQ1)

For a single adult of age i in rating area j at % of poverty in year t , the credit is C_{ijt} , where B_{jt} is the annual benchmark, P_{ijt} the applicable percentage and I_{ijt} income at the state's poverty guideline. The cliff is the difference in net premium between 400% and 401% of poverty, computed rating area by rating area. Statistics across rating areas are medians unless stated.

Cost structure and the pool (RQ2)

Annual expenditure is modeled with a Tweedie GLM (log link, variance power 1.6) on age band, income band, sex, any chronic condition, the count of chronic conditions, region and non-group coverage in any month, with person weights and standard errors clustered on the stratum-PSU pair. The non-group coefficient tests whether the private-market cost structure transfers to the individual market at given observables.

The selection pool is the MEPS sample of privately insured adults aged 18 to 64 at or above 100% of poverty, 50,793 person-years. Their weights are raked to the 2025 age and income marginals of plan selections in the HealthCare.gov states. The raking factors run from 0.03 to 20.9 because MEPS's privately insured sample is concentrated at higher incomes while 54% of marketplace selections are at 100 to 150% of poverty; the effective sample size after raking is about 4,200. Each person's expected cost is the Tweedie prediction; blends with realized cost are reported as sensitivity.

Each person is priced at the median benchmark for their own single year of age in each year, and pays the smaller of that premium and their expected contribution. The price shock is the rise in the person's own payment in percentage points of income, which is defined for the 54% of enrollees who paid nothing in 2025; a proportional metric is reported as sensitivity. The probability of exit is a response parameter times the shock, divided by E_{ijt} , where E_{ijt} is standardized log expected cost and θ is a selection tilt. After exit, the benchmark is scaled by the ratio of stayers' mean expected cost to the baseline mean and exits are recomputed until the factor settles; a factor above 3 is recorded as a spiral. The response is solved by bisection so that simulated retention in the bands at or above 100% of poverty equals the observed 2026 to 2025 ratio in the HealthCare.gov states. A second calibration solves one response for enrollees above 400% of poverty and another below, jointly, so that each segment's retention matches its own observed value.

One market or two (RQ3)

Three tests follow, each with a placebo. First, the observed change in each income band from 2025 to 2026 is set against the change from 2024 to 2025, when the enhanced schedule applied in both years. Second, if the threshold induced households to report income below it, the 300 to 400% band should exceed a counterfactual built from the bands below it; the counterfactual moves the band's starting count at the pooled rate of the 100 to 300% bands, and the same calculation is made for 2024 to

2025. Third, a difference-in-differences regression on state by band by period cells:

where exposed marks the two bands above 400% of poverty, notch marks the 300 to 400% band, post marks 2025 to 2026, are state-by-period effects and band effects. Cells are weighted by the starting count and standard errors clustered by state. A dose version interacts the exposed terms with the state's unsubsidized benchmark at age 60 as a share of income at 401% of poverty, centered, so that the coefficient is per percentage point of burden; a reinsurance version interacts them with Section 1332 status. The same design is run on 2,048 counties with county-by-period effects, using the county's 2025 average premium as the dose. A simpler cross-state regression of the exposed segment's change on burden, weighted by segment size, is reported with its 2024 to 2025 placebo. The repricing loop is also run for representative subsidized and unsubsidized households at each age at the segment-specific calibrated responses and at multiples of them.

Loss ratios (RQ4)

For each year the simple loss ratio is incurred claims over earned premium and the regulatory ratio is the statutory numerator over the statutory denominator. To illustrate the bound the 80% rebate floor imposes in 2026, the 2024 regulatory ratio is carried forward by two years of an assumed claims trend of 4% to 12% and divided by the 2024 to 2026 premium factor, on two premium paths: the average premium of plans selected, and that path with the 2025 to 2026 step replaced by the median benchmark increase.

Cost per person-year of coverage (RQ5)

For each income band the extra credit from a candidate schedule applied to 2026 premiums and poverty levels is computed at the band midpoint, with each age band priced at its

own benchmark and weighted by the 2025 age mix. Coverage gained is the gap between the band's 2025 count moved at the rate of the 100 to 150% band, the control, and its observed 2026 count. The cost is the extra credit paid to everyone the band would then contain; the cost per person-year is that over coverage gained. Above 400% the bunching excess is also removed from coverage gained, since relabeled households were covered throughout. Four schedules are priced: the enhanced schedule in full; the enhanced schedule to 600% of poverty; the enhanced schedule to 400% only; and the 2026 schedule with its 9.96% cap extended above 400%. For partial schedules a band recovers the share of its full-restoration gain equal to the share of the full payment reduction the schedule delivers, an assumption stated as such.

RESULTS

The premium shock (RQ1)

The median benchmark for a 40-year-old rose 24.3% between 2025 and 2026, with an interquartile range across rating areas of 17.5% to 32.8%. Table 1 shows the schedule compounding this. At 150% of poverty a 40-year-old paid nothing in 2025 and \$984 in 2026; at 200% the payment rose 243%; at 400% it rose 22% to 10.0% of income.

At 401% of poverty the credit ends. Figure 1 shows the notch: one more dollar of income costs a 40-year-old \$1,487, where in 2025 it cost \$13. Because the 3:1 age band applies to the gross premium and the percentage cap it replaced did not, the cliff grows with age (Table 2). It is nothing for a median 21-year-old, \$10,163 for a 60-year-old and \$11,891 for a 64-year-old. At 401% of poverty the full benchmark takes 12.3% of income at age 40, 26.1% at 60 and 28.9% at 64, and exceeds a fifth of income in 90% of rating areas at age 60 (Table A1).




Figure 1: Net benchmark premium by income, single adult aged 40, median HealthCare.gov rating area, 2025 and 2026 schedules.




Figure 2: Change in plan selections by income band, 30 HealthCare.gov states, 2025 to 2026 and the 2024 to 2025 placebo.

By state, the median cliff at age 60 exceeds a third of income in West Virginia and Wyoming (Table A3).

Who left

Plan selections fell 4.9% nationally, from 24.32 to 23.13 million. The published platform totals show HealthCare.gov down 7.9% and state-based exchanges up 2.3%, but Illinois's 466,000 selections moved between the two platforms. On a constant set of states HealthCare.gov fell 5.4% and the state-based exchanges 3.9% (Table 3), as the CMS 2026 open enrollment report also notes.

In the 30 HealthCare.gov states the change is dominated by the threshold (Table 3b, Figure 2). The bands that kept a credit fell between 3.0% and 8.2%, except the 300 to 400% band, which grew 8.7%. The 400 to 500% band fell 45.9% and the band above 500% fell 25.8%; together the bands that lost the

credit fell 34.6% against 3.5% for the bands that kept one. In the prior year the same exposed bands grew 9.0%, in line with the rest of the market. By age, selections at 18 to 25 and 26 to 34 fell 8.9% and 10.0% against 6.8% at 55 to 64, a gradient in the direction selection predicts.

The pool (RQ2)

The cost structure transfers

In the Tweedie model (Table 4), non-group coverage has a relativity of 0.943 ($p = 0.468$) with age, income, sex, region and chronic conditions held constant, so the individual market is not distinguishable from the rest of the private market on cost at given observables. Each additional chronic condition multiplies expected cost by 1.41. Mean cost rises 3.12 times from ages 18 to 25 to ages 60 to 64 (Table A11), close to the




Figure 3: Observed change by segment against the selection model with one response, fitted to the total, and with one response per segment.




Figure 4: Rise in the household's own annual payment for the benchmark plan, 2025 to 2026, by age, for three household types.

statutory 3:1 band, so age is priced roughly at cost; within age bands the 90th percentile costs 26 to 69 times the 25th, which is the spread community rating does not price.

One calibrated response

Observed retention in the 100%-and-above bands was 94.8%. The response that reproduces it at a tilt of 1.5 is 0.0143 exits per percentage point of income, with a first-year morbidity rise of 1.2%; across tilts and cost specifications the rise is between -0.1% and 2.6% (Table 5). Pricing everyone at the 40-year-old benchmark, or using the proportional metric, changes the response little. The earlier specification, which post-stratified onto MEPS's own non-group mix and priced everyone at age 40, gave a response four times larger on a different metric and is retained in Table 5 for comparison. No specification approaches the spiral threshold, and the grid in Table A6 shows retention above 83% at responses up to nearly three times the calibrated one.

One response does not fit

The calibrated model predicts a 5.1% loss above 400% of poverty and 5.2% below (Figure 3). The observed losses were 34.6% and 3.5%. Fitting one response to each segment reproduces both, with 0.107 exits per point of income above the threshold and 0.0097 below, a ratio of 11. The ratio is the same in the simpler band-by-band comparison of Table A7, where the observed change per point of income is between -1.3 and -2.0 in the bands below 300% of poverty and -7.6 and -8.5 above 400%. At the calibrated single response, leavers cost 22% less than stayers on average (Table A5); the pool moves little because so few leave.

One market or two (RQ3)

Age decides the size of the shock

For a household at 450% of poverty, the rise in its own payment

for the benchmark is \$1,455 at age 27 and \$10,638 at age 60 (Figure 4, Table A2). A household at 250% sees a larger proportional rise, 119%, from a base of \$1,506. Run separately for each household type at the segment-specific responses, the repricing loop converges everywhere; the premium factor is bounded at for a cumulative exit share when leavers cost 22% less than stayers, and reaches 1.22 for a 64-year-old at 450% of poverty. A spiral in the sense of divergent repricing does not appear in any cell at four times the calibrated responses. What appears instead is a large level shift concentrated among older unsubsidized households.

Bunching below the threshold

Had the 300 to 400% band moved with the 100 to 300% bands, which fell 4.4%, it would have held 898,631 selections in 2026. It held 1,021,383, an excess of 122,752, which is 71% of the 173,158 selections lost in the 400 to 500% band immediately above (Table 6, Figure 5). In the placebo year the same calculation gives a placebo excess of -66,339: the band grew more slowly than its comparison bands when there was no cliff. Carrying that pre-period gap forward raises the 2026 excess to 105% of the loss above. Using the 250 to 300% band alone or the 200 to 300% bands as the comparison gives 84% to 92% before adjustment (Table 6). Aggregate counts cannot show that the same households moved, but growth immediately below a notch, matching the loss immediately above it and absent in the prior year, is what bunching predicts (Kleven and Waseem, 2013). Part of the fall above the threshold is therefore relabeling rather than loss of coverage. A relabeled household draws a credit on projected income it may not have, which is reconciled at tax filing; that cost is not measured here.

Difference in differences

Table 6b reports the state-by-band regressions. With state-by-period and band effects, the exposed-by-post coefficient is -31.8




Figure 5: Plan selections by income band, 30 HealthCare.gov states, 2025 and 2026, with the counterfactual for the 300 to 400% band.

log points (SE 5.0): the bands above 400% fell that much more from 2025 to 2026 than from 2024 to 2025, relative to the other bands, and the 300 to 400% band rose 22.3 log points more. Both survive in the 47-state sample. The dose version places the exposed effect at -3.25 log points per percentage point of unsubsidized burden (SE 0.68); the level term on burden in the placebo period is 0.26 and insignificant, so states with high burdens were not already losing their exposed segments. The county panel gives the same effect, -31.3 log points, with county-by-period effects and a stronger loss in counties with higher 2025 premiums (Table A9).

Across states. Table 7 and Figure 6 show the simple cross-state form. From 2025 to 2026 the slope of the exposed segment's change on burden is -1.369 percentage points per point ($t = -6.57$); from 2024 to 2025 the slope is 0.082 ($p =$

0.59). West Virginia and Wyoming, where the unsubsidized benchmark for a 60-year-old exceeds 43% of income at 401% of poverty, lost 46.8% and 59.7% of the segment (Table A4). Whole-state enrollment shows no gradient, because 95% of enrollees in these states face a capped payment. Part of the slope is plausibly bunching as well as exit, since the incentive to report income below the threshold is larger where the premium above it is larger.

Reinsurance

The seven HealthCare.gov states with Section 1332 reinsurance saw benchmark growth 10.5 points lower than the other 23 (14.7% against 25.2%, $p = 0.036$; Table A8b). Their exposed segments fell less, by 10.8 log points ($p = 0.06$), but the difference shrinks to 4.1 log points and loses significance once burden is held constant (Table A8). Reinsurance




Figure 6: Change in plan selections above 400% FPL against the burden of unsubsidized coverage at age 60, 30 HealthCare.gov states; left, the placebo period; right, 2025 to 2026. Marker area is the segment's size at the start of the period.




Figure 7: Cost of restoring the enhanced schedule per additional person-year of coverage, by income band, 30 HealthCare.gov states, 2026.

appears to work through the level of the premium rather than independently of it, and with seven treated states the test has little power.

Loss ratios and the rebate floor (RQ4)

The individual market's regulatory loss ratio was 90.4% in 2021, 87.8% in 2022, 86.3% in 2023 and 87.1% in 2024 (Table 8), against a statutory floor of 80%, while life-years grew from 14.84 to 22.43 million. Claims per member month rose 1.8% a year over the period, which is partly a composition series: the enrollment the enhanced credits brought in was, on this evidence, cheaper than the enrollment already there. In 2024, 12.0% of life-years were with issuer-state entities below the floor, owing \$1.17 billion of the \$1.23 billion in rebates (Table A10). The median state's Herfindahl index is 0.339 and Alaska, Hawaii and Vermont each had three issuers.

The average premium of plans selected rose 22.5% from 2024 to 2026, with selections shifting from silver to bronze. On that path the 2024 loss ratio carried forward at a 4% claims trend becomes 76.9% in 2026 and the floor returns \$269 per enrollee-year; at a trend above 6.1% it does not bind (Table 8b). On the benchmark path the rebate at a 4% trend is \$536. That bounds the transfer to insurers, not the loss of coverage: a rebate reaches only enrollees, after the reporting year closes. Adverse selection raises claims and therefore the loss ratio, so a pool losing its lower-cost members shows improving compliance with the floor, and the series is a lagging indicator of selection.

The cost of coverage by income band (RQ5)

Restoring the enhanced schedule for 2026 in the 30 HealthCare.gov states would cost \$14.9 billion a year and recover 458,000 person-years of coverage, \$32,525 each, or \$44,435 net of bunching (Table 9). Applied to the 47 states with consistent income reporting the outlay is \$23.5 billion and the average

\$30,786. The average conceals a wide gradient (Figure 7): coverage costs \$9,064 a person-year at 400 to 500% of poverty and \$57,208 at 200 to 250%. In the 150 to 300% bands only 3% to 5% of the extra spending reaches a household whose enrollment it changes, because most of those households stayed; at 400 to 500% the marginal share is 44%. The 300 to 400% band grew, so no coverage gain is attributed to it although it would receive \$1.5 billion.

Table 9b prices the alternatives. Extending the enhanced schedule only to 600% of poverty saves 5% of the outlay and forgoes 13% of the coverage. Restoring it only below 400% costs \$18.9 billion nationally for 319,000 person-years, \$59,059 each, because it spends nothing where the response was largest. Leaving the 2026 schedule in place and extending its 9.96% cap above 400%, which removes the notch and nothing else, costs \$3.3 billion for 338,000 person-years, \$9,636 each, or \$16,672 net of bunching. Under the proportional-recovery assumption that is the cheapest coverage on the menu; it is also the schedule whose recovery assumption is least tested, since no year has combined the original contribution rates with an uncapped credit.

The Congressional Budget Office (2025) estimated that permanent extension would add \$350 billion to deficits over 2026 to 2035 and raise the insured population by 3.8 million in 2035, about \$9,211 per person insured at the end of the window. The full-restoration average here is more than three times that. The two measure different things: a decade-long equilibrium that includes people who never entered the marketplace, against the first-year reversal of one observed change, dampened by the inertia Handel (2013) documents. The gap does not close in the convenient direction, since netting out bunching raises the estimate here.

DISCUSSION

What the aggregate hides

Read at the level of national totals, the 2026 marketplace held: enrollment fell by a twentieth, the calibrated price response is small, and the pool's expected cost moved by about 1%. Each statement averages two segments. Where the schedule caps what a household pays, benchmark linkage absorbed the premium increase and the response to the higher contribution rates was mild. Above 400% of poverty the credit and the insulation ended together, the segment lost a third of its enrollment, and the loss was largest where full-price coverage takes the most income. The placebo year shows none of this, and the difference-in-differences and county estimates agree with the simple comparisons.

Hypothesis H1, that exits concentrate among younger and above-400% enrollees, holds for income and weakly for age. H2 holds in its first half: the subsidized market was stabilized by benchmark linkage. Its second half, spiral dynamics in the unsubsidized segment, appears as a large one-time level shift in older and high-premium cells, with repricing that converges rather than diverges. H3 is tested in part: reinsurance states had lower premium growth and smaller exposed-segment losses, but the enrollment effect is not separable from the premium level with seven treated states. The comparison of restored credits across income bands contradicts the second half of H3, since the coverage recovered per dollar is highest above the threshold.

For regulators and insurers

A regulator monitoring total enrollment, average premiums or loss ratios would see little of this. The movement is in a segment that was 5% of selections and in the boundary between that segment and the one below it. The ratio of selections just below 400% of poverty to selections just above is a direct indicator of the notch's effect and is available in the state-level files. For insurers, the exposed segment is concentrated among older enrollees in high-premium rating areas, where the full premium exceeds a quarter of income; that is where the risk mix can shift fastest and where a pooled rate filing does not show it. The reconciliation of projected against actual income at tax filing will show how much of the bunching was reporting.

Limitations

The planned synthetic population from the American Community Survey with rating-area assignment was not built. The pool is MEPS's privately insured sample raked to marketplace marginals, with an effective sample of about 4,200 and factors above 20 in the 100 to 150% cells; if MEPS's undercount of the marketplace is selective within cells the composition inherits it. Expected cost is a model prediction on observables, so the tilt parameter carries all of the selection on unobserved health. There is no rating-area equilibrium and no simulated loss-ratio distribution; RQ4 is answered from filed experience, and the 2026 loss ratios will

not be filed until 2027. The calibration was not validated on held-out states; the placebo and difference-in-differences tests serve as the falsifiable checks. Enrollment is plan selections, and the effectuation rate may have fallen more for households facing large increases. The premium data exclude state-based exchanges. All calculations are for single adults. Bunching is inferred from aggregate counts, and genuine income changes cannot be separated from reporting. The policy costing is static and its partial-schedule recoveries rest on a proportionality assumption; the 47-state costing applies HealthCare.gov premiums to states whose premiums are not in the data.

CONCLUSION

The expiration of the enhanced premium tax credits did not produce a premium spiral in the 2026 individual market. It produced two segments with different responses. Where the schedule capped payments, enrollment fell by a few percent. Where the credit ended, a third of enrollees left the segment, many apparently for the income band just below it, and the loss was largest in the states where full-price coverage takes most of income. None of this was present in the prior year. A single elasticity describes neither segment, the rebate floor constrains what insurers keep but not who stays covered, and the coverage a dollar of restored credit buys varies more than fivefold across the income distribution. Each of these is visible only when the market is examined in the pieces the subsidy schedule creates.

DECLARATIONS

Data availability

CMS Exchange Rate and Plan Attributes public use files, Marketplace Open Enrollment Period public use files and Medical Loss Ratio public use files are published by the Centers for Medicare & Medicaid Services without registration. MEPS public-use files are available from the Agency for Healthcare Research and Quality. None is redistributed.

Code availability

Complete reproduction code in Python is at <https://github.com/tosin-babs/aca-risk-pool-subsidy-cliff>; python/run_all.py rebuilds every table and figure. An interactive calculator implementing the paper's subsidy calculation is at <https://aca-subsidy-cliff.vercel.app>.

Competing interests

None declared.

Ethics

The analysis uses public administrative files and de-identified survey data and did not require ethical approval.

REFERENCES

1. Akerlof, G. A. (1970). The market for "lemons": quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500. doi:10.2307/1879431

2. Cutler, D. M., & Reber, S. J. (1998). Paying for health insurance: the trade-off between competition and adverse selection. *Quarterly Journal of Economics*, 113(2), 433–466. doi:10.1162/003355398555649
3. Einav, L., & Finkelstein, A. (2011). Selection in insurance markets: theory and empirics in pictures. *Journal of Economic Perspectives*, 25(1), 115–138. doi:10.1257/jep.25.1.115
4. Einav, L., Finkelstein, A., & Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics*, 125(3), 877–921. doi:10.1162/qjec.2010.125.3.877
5. Finkelstein, A., Hendren, N., & Shepard, M. (2019). Subsidizing health insurance for low-income adults: evidence from Massachusetts. *American Economic Review*, 109(4), 1530–1567. doi:10.1257/aer.20171455
6. Geruso, M., & Layton, T. J. (2017). Selection in health insurance markets and its policy remedies. *Journal of Economic Perspectives*, 31(4), 23–50. doi:10.1257/jep.31.4.23
7. Hackmann, M. B., Kolstad, J. T., & Kowalski, A. E. (2015). Adverse selection and an individual mandate: when theory meets practice. *American Economic Review*, 105(3), 1030–1066. doi:10.1257/aer.20130758
8. Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: when nudging hurts. *American Economic Review*, 103(7), 2643–2682. doi:10.1257/aer.103.7.2643
9. Jaffe, S., & Shepard, M. (2020). Price-linked subsidies and imperfect competition in health insurance. *American Economic Journal: Economic Policy*, 12(3), 279–311. doi:10.1257/pol.20180198
10. Kautter, J., Pope, G. C., Ingber, M., Freeman, S., Patterson, L., Cohen, M., & Keenan, P. (2014). The HHS-HCC risk adjustment model for individual and small group markets under the Affordable Care Act. *Medicare & Medicaid Research Review*, 4(3), E1–E46. doi:10.5600/mmrr.004.03.a03
11. Kleven, H. J., & Waseem, M. (2013). Using notches to uncover optimization frictions and structural elasticities: theory and evidence from Pakistan. *Quarterly Journal of Economics*, 128(2), 669–723. doi:10.1093/qje/qjt004
12. Rothschild, M., & Stiglitz, J. (1976). Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. *Quarterly Journal of Economics*, 90(4), 629–649. doi:10.2307/1885326
13. Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2(3), 180–212. doi:10.1257/pol.2.3.180
14. Saltzman, E. (2019). Demand for health insurance: evidence from the California and Washington ACA exchanges. *Journal of Health Economics*, 63, 197–222. doi:10.1016/j.jhealeco.2018.11.004
15. Smyth, G. K., & Jørgensen, B. (2002). Fitting Tweedie's compound Poisson model to insurance claims data: dispersion modelling. *ASTIN Bulletin*, 32(1), 143–157. doi:10.2143/AST.32.1.1020
16. Centers for Medicare & Medicaid Services (2024, 2025). *Health Insurance Exchange Public Use Files: Rate and Plan Attributes*, plan years 2025 and 2026.
17. Centers for Medicare & Medicaid Services (2024–2026). *Marketplace Open Enrollment Period Public Use Files, State-Level and County-Level*, 2024, 2025 and 2026, and *Health Insurance Exchanges 2026 Open Enrollment Report*.
18. Centers for Medicare & Medicaid Services (2022–2025). *Medical Loss Ratio Public Use Files*, reporting years 2021–2024.
19. Agbonome, T. N. (2026). Influence of dietary choices on maternal health during pregnancy among individuals with PCOS in low-income communities. *Current Medical and Drug Research*, 10(02), 14–22.
20. Centers for Medicare & Medicaid Services (2022, n.d.). *Data Brief on Section 1332 Waivers: State-Based Reinsurance Programs* (December 2022) and *Section 1332: State Innovation Waivers* (web page).
21. Agency for Healthcare Research and Quality. *Medical Expenditure Panel Survey Full-Year Consolidated Files*, 2019–2024.
22. Agbonome, T. N. (2026). Management of mpox in Africa: Ongoing threat and epidemiological insights. *Current Medical and Drug Research*, 10(02), 1–13.
23. Internal Revenue Service (2025). *Revenue Procedure 2025-25: Applicable Percentage Table and Required Contribution Percentage for 2026*. Washington, DC.
24. Office of the Assistant Secretary for Planning and Evaluation (2024, 2025). *HHS Poverty Guidelines*. Washington, DC: U.S. Department of Health and Human Services.
25. Congressional Budget Office (2025). *The Estimated Effects of Enacting Selected Health Coverage Policies on the Federal Budget and on the Number of People With Health Insurance*. September 2025. <https://www.cbo.gov/publication/61734>