

Deadpooling: Strategic Lock-In in the Medigap Market

UC Berkeley

Margaret Kallus* and Elaine Shen[†]

October 25, 2026

[Click here for the latest version](#)

Abstract

The Medigap market prevents adverse selection by locking enrollees in to their plan; after an initial enrollment window, insurers may reject applicants who fail medical underwriting, so most enrollees have only one opportunity to shop. We show that insurers profit from this lock-in through corporate structure. Linking two decades of Medigap filings to parent–subsidiary ownership records, we document *deadpooling*. In this practice, one subsidiary closes a plan to new enrollees, a sister subsidiary opens a replacement for healthy applicants, and premiums rise on the enrollees left behind. From 2006–2020, 17.6% of Medigap life-years are spent in plans that deadpool. Using a Callaway–Sant’Anna staggered difference-in-differences design, we find that an enrollee who stays in their plan pays \$762 more in premiums five years after closure than if their plan had not closed. Each year’s closures transfer \$64.7 million from seniors to insurer profits and cost seniors \$108.3 million in lost consumer surplus, including switching costs and lost coverage. Overall, seniors lose \$1.67 in consumer surplus per dollar of insurer gain. With variation from the staggered adoption of switching laws in six states, we find that allowing enrollees to switch plans annually reduces enrollment in deadpooling plans by 18%. A calibrated model shows that this can be explained by consumer inattention, and simulates market outcomes under alternative policies.

*UC Berkeley Economics PhD Program, margaret_kallus@berkeley.edu

[†]UC Berkeley Economics PhD Program, elaineshen@berkeley.edu.

1 Introduction

A well-functioning insurance market protects consumers from adverse financial consequences when they experience negative events. However, these markets are vulnerable to adverse selection, which occurs when buyers know more about their own risk than insurers are able to price. In health insurance markets, the inability to include some information in prices often comes from legislative mandates, in which insurers cannot deny coverage to people with preexisting conditions or health risks during certain enrollment periods. Left unchecked, adverse selection can cause a market to unravel. This can leave a competitive market with no stable equilibrium at all (Rothschild and Stiglitz 1976), and it has been documented across a wide range of insurance settings (Einav and Finkelstein 2011, 2023). To prevent unraveling, regulators use a variety of tools to encourage low-risk enrollees to purchase coverage, including subsidies and mandates. However, subsidies may be expensive to administer, and mandates are often unpopular, especially in US health insurance markets. Another way to encourage low-risk beneficiaries to purchase insurance coverage is by asking consumers to make long-term choices. Risk-averse enrollees would be motivated to purchase a higher level of coverage for a long-term plan than they would for a short-term plan, as their long-term risk is more uncertain. However, if there are barriers to switching between insurers, long-term contracts can give firms ex-post market power over their enrollees.

We study the costs of using long-term choices in the market for Medicare Supplement insurance, or Medigap. Medicare, the US's federal health insurance program for senior and disabled populations, is particularly susceptible to threats from adverse selection, as these groups are especially likely to require costly care. Medicare expenditures account for 3.8% of US GDP, much of which is subsidized by the government, so problems in this market carry budgetary and macroeconomic repercussions (Centers for Medicare and Medicaid Services 2025). The Medicare program places bounds on how sparse or comprehensive approved coverage may be, but within these bounds individuals choose among a wide range of plans differing in benefits and cost-sharing structures. Medigap policies sit at the comprehensive end of that range, covering most cost sharing without restricting access to care.

Because Medigap coverage is high-quality but expensive, one might expect a regime in which enrollees could freely opt in and out of this coverage to suffer from adverse selection. In practice, seniors would wait until they developed health conditions to pay for a Medigap and enroll in inexpensive Medicare Advantage plans while their utilization was low. Curto (2023) finds exactly this in the small number of states with highly flexible plan switching regulation; when people can buy a Medigap at any time they like, they delay purchase, which is associated with decreased enrollment and higher annual premiums.

To prevent these effects, most states require beneficiaries to choose whether to get a Medigap plan when they first join Medicare. A beneficiary is guaranteed access to any Medigap plan when they first enroll at 65, but after this initial period they must pass medical underwriting to switch insurers or policies.

Anyone who cannot pass medical underwriting is locked in to their initial plan choice. In this paper, we find that some firms leverage this lock-in for profit. Insurers use subsidiaries to segment enrollees by health status, charging a captive pool above cost while competing for healthy enrollees under a different

corporate name. We call this practice *deadpooling*.

This general mechanism is not peculiar to the Medigap market. A seller who can distinguish locked-in buyers from still-mobile ones, and price to them separately, holds ex-post market power that no degree of ex-ante competition disciplines. This dynamic could appear in any long-term relationship in which a seller can reset the price over time and the buyer's ability to walk away erodes, including the owner of equipment whose manufacturer uses proprietary parts (Borenstein, MacKie-Mason, and Netz 1995) or an aging renter who can no longer easily pack up and move.

Though this dynamic appears in other settings, the Medigap market is a particularly good place to study it, as it has standardized products. This makes it possible to disentangle changes in price from changes in plan quality and benefits. Because any two plans carrying the same letter provide identical coverage, a premium increase after closure is a pure financial loss. We construct a novel Medigap pricing and plan choice dataset from insurance filings with the NAIC (National Association of Insurance Commissioners), annual parent–subsidiary linkages from the NAIC, and highly granular pricing data from CSG Actuarial in order to estimate the prevalence and impact of deadpooling.

This paper makes three main contributions. First, we document a new way that firms exploit the market power lock-in creates. Work on adverse selection typically views long-term contracts and limits on plan switching as a remedy, and measures their costs in reclassification risk or forgone choice (Handel, Hendel, and Whinston 2015; Atal 2019). We show a further cost: because the regulations that discipline lock-in typically bind at the subsidiary level rather than the parent level, insurers can sort enrollees by health across corporate entities, raise prices on the captive pool, and keep competing for the healthy. This differs from the frictions studied in the switching-cost literature (Klemperer 1987, 1995; Farrell and Klemperer 2007), as an ordinary switching cost is surmountable when faced by an attentive consumer. Here, the barrier is prohibitive and falls on the buyers with the highest health risks and fewest outside options. We find that 17.6% of Medigap life-years are spent in plans that deadpool.¹

Second, we measure what this costs consumers, separating the part of the premium increase that reflects the selected, higher-risk pool from the part that is additional markup. To do so, we use quoted prices from CSG Actuarial, which are new to this literature; they record what an insurer charges a given applicant, so unlike realized premiums they do not move with the composition of the pool. Using a Callaway–Sant'Anna difference-in-differences design to account for the staggered timing of plan closures and heterogeneous treatment effects, we estimate that each year's closures transfer \$64.7 million in excess markup from seniors to insurer profits over the first five years. This totals about \$1 billion over our sample period. A senior who stays in a deadpooled plan faces increasingly higher quoted premiums each year after closure, rising from \$151 to \$762 over the first five years. These price increases do not buy any additional changes in insurance benefits, and the beneficiary would not face them had they chosen a different insurer.

The magnitude of harm caused by deadpooling is large compared to other estimates in the literature.

1. This share counts a deadpooling plan's entire life, including the years before it closes, and so measures exposure to a plan that deadpools rather than time spent trapped. The narrower share of life-years spent inside a block that has already closed, which is the window over which harm is computed, is 9.4%.

Ericson (2014) attributes a markup of about 10%, roughly \$50 per enrollee per year, to consumer inertia in Medicare Part D, and Atal (2019) estimates that enrollees in Chile’s guaranteed-renewable market would pay 13% more for the ability to switch. In our setting, quoted premiums have risen by 27% by the fifth year after plan closure. Including switching costs and the consumer surplus given up by seniors who are priced out by deadpooling, each year’s closures cost enrollees \$108.3 million in consumer surplus. Of this, \$64.7 million is a transfer from seniors to insurer profits, \$43.0 million is lost surplus from leaving Medigap, and \$0.7 million is switching costs. Put another way, each dollar of profit insurers gain from deadpooling costs seniors \$1.67.

For our third contribution, we provide the first evaluation of the state switching rules passed to curb this practice. A fully attentive consumer could use switching rules to leave a deadpooling plan, which would reduce the theoretical harm of this practice to almost zero. In practice, these rules reduce enrollment in deadpooling plans by only 18%. A calibrated model attributes this gap to inattention. We also consider alternative policies, such as banning deadpooling and combining switching rules with automatic switching by a regulator (as in Part D’s Low-Income Subsidy program), and find that automatic switching is more effective than policies that rely on attentive consumers to improve market outcomes.

The rest of the paper proceeds as follows. Section 2 describes the Medigap market, its regulation, and the mechanics of deadpooling, and Section 3 describes related literature. Section 4 discusses the data and Section 5 our empirical strategy. Section 6 documents how common deadpooling is, what happens to premiums, claims and profits after a block closes, and whether the practice is profitable and chosen. Section 7 quantifies the cost to consumers. Section 8 evaluates state switching rules, and Section 9 uses the model to compare them with alternative reforms. Section 10 concludes.

2 Background on the Medicare Market

2.1 Medicare and Medigap

The Medicare market insures over 65 million seniors and individuals with disabilities. Upon enrolling in Medicare, beneficiaries decide whether to enroll in Original Medicare (also called Fee-For-Service or Traditional Medicare) under the federal government or a Medicare Advantage (MA) plan from a private company. Original Medicare has essentially no restrictions on which doctors you can see, and it does not require prior authorization to cover procedures or see specialists. However, coverage is limited; there is a fairly large hospital deductible, and it only covers 80% of the cost of most appointments and procedures. This leaves beneficiaries responsible for the remaining 20% coinsurance, with no out-of-pocket maximum to cap annual costs.

Medicare Advantage plans are the alternative to Original Medicare. These plans are offered through private insurers and typically have more restrictions on care through networks, required referrals for specialist care, and prior authorization before some appointments and procedures. MA plans also place limits on cost sharing with annual out-of-pocket maxima, and they often offer supplemental benefits like dental, vision, and hearing.

Most beneficiaries with Original Medicare have supplemental coverage to fill the gaps in Original Medicare's cost sharing. Some beneficiaries who fall below income and asset thresholds are dually eligible for Medicaid, which covers the cost sharing directly (about 18% of enrollees).² Others purchase a Medigap plan. These plans are sold by private insurers and offer some of the most comprehensive coverage available, substantially reducing or even eliminating altogether out-of-pocket cost when accessing medical care. About 14 million people in the US today have a Medigap.

Since the passage of the 1990 Omnibus Budget Reconciliation Act (OBRA), Medicare Supplement plans are standardized from a set menu of plan types (Freed et al. 2023).³ All standard options are fairly comprehensive, either capping out-of-pocket annual spending or covering costs almost entirely. The full benefits provided by each plan type are shown in Appendix Figure A1.⁴ Medigap insurers are secondary payers, so they must cover cost sharing if Original Medicare approves the expense, and they cannot cover cost sharing if Original Medicare denies their portion of the expense. This leaves very little discretion left to the insurer on what expenses to cover, which reduces the administrative burden of offering these plans. Because Original Medicare is responsible for claims approval, Medigaps function largely as financial products, limiting the scope for differences in customer service or other firm characteristics to impact their customers for this specific product.⁵

Medigap plans can carry high premiums, particularly when contrasted against MA. We calculate that average paid premiums for Medigaps exceed \$200 per month, while three quarters of enrollees in individual Medicare Advantage drug plans pay no premium beyond the Medicare Part B premium (Freed et al. 2026). Despite these high premiums, a focus group of insurance agents conducted by the Commonwealth Fund reports that most agents plan to enroll in Medigap themselves upon retirement, suggesting that expert consumers value this comprehensive and flexible coverage even at substantial monetary cost (Ali et al. 2021).

2.2 Lock-In: Guaranteed Issue and Medical Underwriting

Because Medigap is the most comprehensive insurance product available in this market, adverse selection is of particular concern. To address this, federal regulations limit the periods when Medicare beneficiaries can get Medigap coverage. Seniors are eligible to purchase any Medigap plan at the lowest available rate for the 6 months after they turn 65 (or first get Part B). If they miss this enrollment window, or they want to switch their plan outside of it, most states require them to undergo medical

2. Authors' calculation from the 2023 Medicare Current Beneficiary Survey public use file, using the survey weights (Centers for Medicare & Medicaid Services 2023).

3. Standardization admits one narrow exception. Section 9.1 of the NAIC Model Regulation permits an insurer to attach "new or innovative benefits" to an otherwise standardized plan, provided they do not alter or reduce any standardized benefit and do not cover outpatient prescription drugs. In practice, these are wellness extras sold alongside the plan, most commonly a gym or fitness membership (such as SilverSneakers) (National Association of Insurance Commissioners 2025). The NAIC filings identify these plans, and we observe that they are uncommon; less than 0.9% of life-years carry an innovative benefit.

4. This is in contrast to markets such as Medicare Advantage, where unusual plan benefits require assumptions. For example, beginning in 2020 some of Anthem's Medicare Advantage plans offered chronically ill members a home pest control benefit and support for members who rely on a service dog (Anthem Blue Cross and Blue Shield 2020). Ranking such a plan against a competitor requires assumptions on the value of some nonstandard benefits. Medigaps pose no such problem: within a plan letter, the benefit is identical across insurers.

5. Interviews with multiple insurance agents have confirmed that consumers seldom contact their Medigap insurer in practice. The main exception would be the limited coverage of health events while traveling internationally offered by some plan types.

underwriting. In practice, this requires enrollees to provide detailed health histories, diagnoses, and past or planned medical services. Insurers may also ask seniors to see a doctor for a recent health screening. Medical underwriting is often binding; we estimate that 73.2% of Medigap enrollees have health conditions that would cause them to fail underwriting (Appendix B). For these consumers, their initial insurer has an effective monopoly over their business, because they would not be approved to buy plans from other insurers.

After their initial guaranteed issue period, there are limited federal consumer protections for Medigaps. Seniors may gain access to special enrollment periods after events such as moving out of a plan's service area, their insurer terminating its contract with CMS, or proving they were misled by an insurance agent or insurance company employee. Changes to a Medicare Advantage plan's provider network, benefits, or prices do not trigger Medigap special enrollment periods, nor do changes to a beneficiary's health status or utilization. While seniors are only guaranteed access to Medigap plans at particular times, each fall any Medicare beneficiary can switch their coverage to any MA plan sold in their area, creating a "one-way door" to MA coverage.

2.3 Rules on Insurers, and Where They Bind

Aware of the market power that lock-in creates, regulators restrict how insurers can act in this market. Insurers set prices for each product at the sex $\times$ ZIP code $\times$ consumer age $\times$ tobacco status level, but they cannot price in health conditions during guaranteed issue periods. These restrictions on pricing could help to spread risk among high- and low-cost enrollees. Insurers who withdraw from the market cannot sell Medigaps again for five years, to induce them to remain in the market. Insurers must also abide by a minimum loss ratio, spending at least 60% of what they collect in premiums on claims in the individual market.⁶ The loss ratio is intended to keep premiums down by limiting allowable premiums relative to costs (Frank and Milhaupt 2022), though in this market it is not typically binding in practice.⁷ States also limit the number of separate risk pools a company can have in the market at a given time to avoid sequestering high-risk seniors in expensive plans. Most importantly, to raise premiums insurers must submit their proposed rates to each state with a justification. Because each of these rules applies to the subsidiary that files the plan, a parent company with several licensed subsidiaries can satisfy every rule within each subsidiary while coordinating prices and plan openings across them.

2.4 Deadpooling

Interviews with insurance agents and actuaries specializing in this market report that large insurance companies have found a way to segment the market while complying with state regulations through a practice known as "deadpooling." This practice relies on high risk individuals being locked in to the first

6. The corresponding floor is 65% in the group market; we work primarily with the individual market throughout this paper. An 85% loss ratio is enforced in many other health markets, such as Medicare Advantage and large-group commercial coverage under the Affordable Care Act (42 CFR 422.2410; 42 U.S.C. 300gg-18).

7. Loss ratio regulations are enforced based on an actuary's 20-year projection of future premiums and claims rather than on a particular year's realized loss ratio, which makes loss ratio violations unlikely to be enforced (NAIC 2000).

plan they buy, as they are unlikely to pass the medical underwriting required to switch.⁸ Here, we describe a simple example of deadpooling. First, an insurance company enters the Medigap market under Subsidiary A, and some seniors join a Medigap plan operated by that subsidiary. Second, Subsidiary A stops accepting additional beneficiaries. Then, the parent company has Subsidiary B enter the market. Healthy and young Subsidiary A beneficiaries with low expected medical costs can undergo medical underwriting and switch to the new, cheaper plan. However, beneficiaries with health conditions are unlikely to pass medical underwriting. This denies them access to a reduced premium. As these higher-cost seniors become concentrated in Subsidiary A, the insurer cites higher healthcare costs as justification for further premium increases. If a sick beneficiary is unable to afford their new, higher premiums, they are priced out of the Medigap market entirely and lose access to this high-quality insurance product.

Deadpooling therefore requires two coordinated actions: one subsidiary closes its plan to new members, and another subsidiary of the same parent opens a new plan in the same market. By separating high- and low-cost beneficiaries into different plans, parent companies can raise premiums on high-cost enrollees without losing market share among low-cost enrollees. Beneficiaries in a plan whose parent company reenters within five years of their plan's closure experience deadpooling. We define this reentry period to be within five years in accordance with the state regulations that prohibit reentry into the market within five years. Figure 1 summarizes the cycle.

2.5 State Responses: Switching and Open-Enrollment Laws


Since 2007, sixteen states have passed additional legislation intended to increase consumer choice within the Medigap market. Here, we define switching legislation as a rule allowing beneficiaries to switch out of their plan and get another Medigap policy at least annually. Figure 2 displays the states that have passed switching legislation. The six states in green (California (2009), Oregon (2013), Nevada (2022), Idaho (2022), Illinois (2022), and Louisiana (2022)) passed switching legislation during our period of study, where we can observe at least three years before and after the law took effect. The remaining states passed their legislation outside our period of observation.

The three states in yellow have even more generous enrollment protections. Open enrollment states require at least annual opportunities for people with no Medigap plan to purchase one without underwriting. This legislation provides more expansive access than switching legislation, but also may be more susceptible to adverse selection and delayed purchase of coverage.

In a classical framework, these regulations would make switching possible for all enrollees. This should increase enrollment in less-expensive plans while decreasing enrollment in deadpooling plans. If perfectly rational enrollees trapped by underwriting could leave at their birthday or plan anniversary, a parent company would gain little from closing a block, since the enrollees it used to keep at a higher


8. Deadpooling is related to, but distinct from, the classic "invest-then-harvest" pattern, in which an insurer prices a cohort below cost early to attract enrollees and then raises premiums once they are locked in. Invest-then-harvest describes a price path within a single risk pool, and attentive enrollees can avoid overpaying by switching from the harvesting plan. Deadpooling requires multiple simultaneous pools, and attentive enrollees with health conditions cannot avoid overpaying.

Figure 1. The Deadpooling Cycle



Notes: The figure illustrates the deadpooling mechanism; the dots track a stylized pool of enrollees, green for low-risk and red for high-risk, and the dollar signs indicate the premium charged on the block. A parent opens a subsidiary block and attracts a mix of low-risk and high-risk enrollees (step 1), then closes the block to new members and raises premiums on the now-captive pool (step 2). Low-risk enrollees who can pass medical underwriting leave for cheaper open plans, while high-risk enrollees who are denied underwriting cannot leave. They remain in the closed block and pay the rising premium or drop Medigap coverage altogether. The parent then opens a fresh low-priced block in the same market through a *different* subsidiary (step 3), which is not subject to the closed block's reentry restriction, and the cycle repeats.

Figure 2. Sixteen States Have Adopted Switching or Open-Enrollment Protections



Notes: Each state is classified by its Medigap switching/open-enrollment regime and the period of adoption. The six states shaded green adopted switching legislation inside our observation window and identify the event studies of Section 8.2. The yellow open-enrollment states and the remaining protected states adopted before the panel begins, so they are excluded from the control group rather than used as treated units. Massachusetts, Wisconsin, and Minnesota (shaded separately) are federal-waiver states that standardize Medigap under their own benefit designs rather than the lettered federal plans, and they are excluded from the sample entirely.

price could follow the healthy into the new subsidiary's plan. However, by making it easier for high-cost enrollees to stay in the market, this legislation could drive up overall premiums. A switching rule is narrower than the continuous open enrollment that Curto (2023) studies; it lets current Medigap enrollees change plans, but it does not let beneficiaries who dropped Medigap buy it back, which limits the scope for waiting to buy coverage until health declines.

A more behavioral framework may suggest that offering consumers additional choice in a market where they make costly mistakes could allow for full or partial correction of these mistakes. However, if enrollees in deadpooled plans are inattentive, a switching rule would help only those who use it. Whether a switching rule protects trapped enrollees is therefore an empirical question which, to our knowledge, has never been studied.

3 Related Literature

This paper contributes to three strands of research: consumer choice and mistakes in complex settings, dynamic insurance markets with lock-in and open-ended pricing, and empirical research on the Medigap market.

The first strand studies mistakes in consumer choice over complex health and financial decisions and the firms that profit from them. Consumers make dominated choices in 401(k) decisions (Choi, Laibson, and Madrian 2011), investment decisions (Mullainathan, Nöth, and Schoar 2011), health decisions (Baicker, Mullainathan, and Schwartzstein 2015), and insurance decisions (Bhargava, Loewenstein, and Sydnor 2017; Handel 2013), and such choices are especially prevalent among senior citizens in the highly complex Medicare market (Handel and Kolstad 2015; Handel and Schwartzstein 2018; Abaluck and Gruber 2011, 2016; Ho, Hogan, and Scott Morton 2017; Ketcham, Lucarelli, and Powers 2015). This literature largely attributes mistakes to consumers, through inattention, limited information, or bounded rationality. On the supply side, firms can shroud add-on prices (Gabaix and Laibson 2006), steer inattentive consumers toward more expensive options (Heidhues, Köster, and Kőszegi 2023), and raise prices on existing customers when switching is costly (Farrell and Klemperer 2007). Competition does not eliminate the rents from these practices. In the shrouding literature, the rent comes from consumers who do not pay attention to prices, and solutions such as disclosure and cheaper switching target attention. In Medigap, most consumers cannot pass underwriting to leave, so more attention helps only enrollees who can still switch. Closest to our setting, insurers in Medicare Part D raise prices on inattentive existing enrollees while launching cheaper new plans (Ericson 2014). Deadpooling is a supply-side strategy of this kind, but it works through corporate structure instead of product design, and it adds health-based lock-in to inattention. This makes choices particularly difficult to reverse and sustains larger markups.

The second strand studies dynamic insurance markets in which enrollees become locked in as their health declines. In this literature, adverse selection is often understood as a threat to the existence of a market: it can unravel a pool (Cutler and Reber 1998) or arise after enrollment from the timing of claims

(Cabral 2017). Restricting when consumers can buy or switch coverage is one way to prevent it, and the costs of this are usually measured as reclassification risk or forgone choice (Handel, Hendel, and Whinston 2015; Atal 2019). We document a different cost of solving adverse selection with lock-in. In deadpooling, insurers engineer death spirals contained to individual blocks in order to raise their margins during those spirals. The spiral then becomes a source of profit rather than a threat to the market. Long-term contracts can protect enrollees by front-loading premiums in exchange for a committed price path (Hendel and Lizzeri 2003; Ghili et al. 2024). In our setting, underwriting commits the enrollee to the plan, but the insurer can still raise the block's premium every year.

The Medigap market is heavily regulated, and a small literature studies how the market works under these rules. Finkelstein (2004) finds that imposing binding minimum standards on Medigap policies was associated with a 25% decline in coverage, consistent with adverse selection, and Cabral and Mahoney (2019) show that Medigap coverage raises an enrollee's Medicare spending by 22%. Starc (2014) treats each parent as one firm and finds that firms hold market power among new enrollees despite standardized products. We look inside the parent, at how it divides existing business across subsidiaries.⁹ Fang, Keane, and Silverman (2008) document that aggregate selection into Medigap is advantageous, but we show that insurers can create adverse selection within a block by sorting enrollees across subsidiaries. Curto (2023) finds that open enrollment and community-rating regulations (which respectively allow all Medicare enrollees to get a Medigap at any time and bar pricing on age, sex, and health) delay enrollment, induce adverse selection, and raise premiums. These papers look at older periods in this market, before Medicare Advantage reached its current market share, while our panel runs through 2025. Lin and Wildenbeest (2020) estimate search costs in choosing a Medigap plan nonparametrically and put the median at \$30, which suggests that the choice of dominated Medigap policies is driven by information frictions more than by the cost of looking.¹⁰ Low search costs do not protect the enrollees targeted by deadpooling; a shopper with zero search costs still cannot act on a better price if medical underwriting bars them from switching. Overall, these papers study pricing and selection when enrollees first buy coverage. We instead follow enrollees after they are locked in, tracing how premiums on a closed block evolve over the years after closure. We also introduce novel data to this setting with quoted premiums from CSG. This lets us understand the options consumers have when enrolling and allows us to separate the price an insurer charges from the realized, paid premium per enrollee, which moves with the composition of the pool. Finally, we provide, to our knowledge, the first evaluation of the state switching regulations meant to address deadpooling.

9. In this respect deadpooling resembles the regulatory arbitrage Kojen and Yogo (2016) document among life insurers, which cede liabilities to reinsurers they own in more lightly regulated jurisdictions. In both cases each legal entity complies with a rule that the group as a whole avoids, but in deadpooling the cost falls on enrollees rather than on the regulator or the public balance sheet.

10. The working paper Maestas, Schroeder, and Goldman (2009) estimates a higher average search cost of \$72 in the same market. To be conservative, we use the lower figure throughout.

4 Data

We construct a Medigap pricing and plan choice dataset from 2006–2025 by combining three main data sources: (i) plan enrollment, premium, and claims filings from the NAIC; (ii) annual parent-to-subsidiary ownership mappings, also from the NAIC; (iii) quoted-rate data from CSG Actuarial, the private consulting firm that maintains the pricing data underlying the government’s Medicare.gov plan finder. The NAIC filings and ownership mappings let us document deadpooling, a firm and provide the first statistics on how common it is. CSG rates let us separate quoted prices from realized premiums, which are affected by selection into coverage. CSG also provides data on rate changes after plans close and the medical conditions insurers use for underwriting.

We supplement this analysis with other data sources including a custom Medigap shopping-behavior module we fielded on the Understanding America Study (UAS), a survey of agents and brokers who enroll seniors in Medicare products, and information from the Medicare Current Beneficiary Survey (MCBS). The UAS module measures how actively enrollees shop, which disciplines the demand side of our model. We get supplemental and contextual information from the Medicare Current Beneficiary Survey and a survey of agents and brokers who enroll seniors in Medicare products.

4.1 NAIC Plan Data

The NAIC insurance filings contain information about the number of beneficiaries covered by each plan at the insurer $\times$ state $\times$ standardized plan type $\times$ group status $\times$ year level from 2006–2025.¹¹ These data were purchased from Mark Farrah Associates. They also contain information on total premiums collected and total claims paid, and some plan characteristics. This dataset is where we obtain the dates when plans were first approved to join the market and when plans were closed or withdrawn from the market.¹² In this analysis, we focus on standardized plans in the individual market.¹³

We link plans across years using the NAIC code of the issuing company, the state, the standardized plan type, and the group status of the plan. The unit of observation in the resulting panel is a *block* (a subsidiary $\times$ state $\times$ standardized plan type cell), which is observed annually. Additional details on cleaning this dataset are reported in Appendix A.2.

From these data, we construct variables used throughout the paper. The financial variables are

11. Covered lives are counted as of December 31 of each reporting year (National Association of Insurance Commissioners 2022, pp. 671–672). Changes made during Medicare’s annual open enrollment period (October 15 to December 7) take effect on January 1 (Centers for Medicare & Medicaid Services 2026), so much of the switching between plans happens at the turn of the calendar year. This makes December 31 numbers more accurate than if plan changes happened continuously throughout the year.

12. Plan closure refers to a pool that no longer accepts new members, but continues to insure existing members who pay their premiums. This does not trigger a special enrollment period for members. Plan withdrawal refers to an insurer leaving a market entirely. Withdrawn plans do not continue to insure existing customers in that market, and plan withdrawal would trigger a special enrollment period for beneficiaries losing their coverage.

13. We exclude employer group plans, as no particular employer plan is in the choice set of the average enrollee. Some retirees have supplemental coverage subsidized by a former employer or union. In the 2023 Medicare Current Beneficiary Survey, 37% of beneficiaries in Original Medicare have employer-sponsored coverage, compared with 47% who have individually purchased private coverage such as Medigap (Centers for Medicare & Medicaid Services 2023). The one group plan we keep is sold by the largest Medigap insurer through a national membership organization. It is filed as a group policy, with the organization as the policyholder, but anyone who joins the organization can buy it, so it is open to the public in the same way an individual plan is. That insurer’s group plans hold 23% of covered lives in the NAIC filings over 2006–2025, and the other group plans we drop hold 2.4%. We also exclude Medicare SELECT plans, as they incorporate provider network restrictions.

per-enrollee premium (total premiums collected divided by covered lives), per-enrollee claims (total claims paid divided by covered lives), per-enrollee profit (total premiums minus total claims, divided by covered lives), and the loss ratio (total claims divided by total premiums; the share of premiums spent on medical costs). All dollar figures are inflation adjusted to January 2025 dollars using the January value of the Consumer Price Index for All Urban Consumers (CPI-U, all items, seasonally adjusted) for each year. Finally, we construct our key object, a *deadpooling* flag, equal to one if the block ever closes and the parent opens a successor subsidiary in the same market (standardized plan type $\times$ state) within five years of closure. Under this definition, 17.6% of life-years through 2020 are spent in plans that deadpool,¹⁴ and about 38% of enrollees in plans that close see their parent reenter the market within five years. The share is similar with a one-, three- or seven-year reentry window (13.3%, 16.2% and 21.0%; Table A11); Appendix D.2 gives the timing of reentry, and Appendices D.3 and D.4 show that the practice does not depend on successors carrying unfamiliar names or on parents steering priced-out enrollees into their own Medicare Advantage plans.

4.2 Subsidiary Mapping Data

We use parent to subsidiary mappings from the NAIC for each year from 2006 to 2025. We link parents to subsidiary codes separately each year in order to ensure smaller companies that change parent ownership over our period are linked to their correct parents. NAIC codes that do not merge to a parent are considered to be their own parent. Closure and reentry are regulated at the subsidiary level, so a parent with many subsidiaries can close one block and open another without violating the regulation. Appendix A.2 describes which edition of the mapping we use and how NAIC codes are matched to years.

4.3 CSG Pricing Data

We supplement the NAIC filings with quoted-rate data from CSG Actuarial, the private firm that maintains the pricing information underlying the government’s Medicare.gov plan finder. We use these data from 2014 to 2025. We use these data at the NAIC subsidiary $\times$ ZIP code $\times$ standardized plan type $\times$ enrollee age level. These data are distinct from the NAIC data on premiums collected, as CSG records the quoted price a new applicant would be offered, whereas the NAIC records the average premium paid by the existing pool who decided to purchase coverage. In a pool that has older enrollees, the average paid premium in the NAIC would be higher than the average expected premium from CSG. From these data we construct quoted rates at different ages. Appendix A.2 contains additional details on how this dataset is cleaned.

A separate CSG file records the rate changes carriers file on blocks that have already closed. Each row is at the NAIC subsidiary $\times$ state $\times$ effective date $\times$ standardized plan type level, and each increase changes all enrollees’ rates proportionally. Because these are the prices charged to enrollees who can no

14. Both the numerator and the denominator of this share are restricted to 2020 and earlier, because identifying a deadpooling block requires observing whether the parent reopens within five years; with the panel ending in 2025, a block closing after 2020 cannot yet be classified.

longer be quoted a new rate, they measure what happens premiums as the pool becomes more selected. Not all states require companies to file rate increases in this way, so these closed-rate data cover 38 states.

4.4 UAS Shopping Module

To measure how actively Medigap enrollees shop, we fielded a custom module on the Understanding America Study (UAS), a nationally representative internet panel run by the University of Southern California's Dornsife Center for Economic and Social Research. Our module (fielded July 2025) asks Medicare beneficiaries about their Medigap enrollment, the number of insurance agents they consulted, and whether they have ever switched plans. It is linked to the panel's standard demographic file.

4.5 MCBS and Medigap Underwriting Criteria

Whether a senior can leave a closed block depends on whether they would pass medical underwriting. We build our estimate of underwriting approval from two sources. The first is the public use file of the Medicare Current Beneficiary Survey (MCBS), a nationally representative survey of Medicare beneficiaries reporting diagnosed conditions, coverage type, and spending. We use these data both for condition prevalence and for the age gradients in claims. The second is the set of underwriting criteria carriers apply, which we collect from CSG Actuarial. These data give the conditions and lookback windows on which an application is flagged for medical underwriting, triggering higher rates or denial of coverage. We match health conditions among Medicare beneficiaries from the MCBS to the conditions used for underwriting in the CSG underwriting criteria data to estimate the share of a closing block's enrollees who would be approved elsewhere at competitive rates. Appendix B contains additional details on the CSG underwriting data and the crosswalk to the MCBS conditions. [xx c: our independent estimate aligns closely with recent work on this topic from KFF—please add a cite and this sentence where appropriate, either here or in the text later on.]

4.6 Medicare Agent Survey

Agents are the intermediaries who place seniors into plans and who see the closed blocks their clients sit on. To ask agents how they advise clients whose block has closed and what they know about the practice of deadpooling, we fielded a survey of licensed Medicare insurance agents. We report some descriptive results in Appendix F.

4.7 Summary Statistics

Table 1 reports summary statistics for our panel. Column 1 has statistics for the full market. Column 2 reports the same statistics for the subset of plans that never close, our main control group. Columns 3 and 4 report statistics for deadpooling plans before and after closure. The top panel displays quoted pricing information, and the bottom panel displays information from company filings with realized financial information and enrollment. We observe 228 parent companies in our sample, which operate

17,101 risk pools. 128 of these parent companies run a block that meets the definition of deadpooling during our sample period.

In the CSG panel, quoted rates for open deadpooling plans and plans that never close are similar, particularly at earlier ages when most enrollees select their coverage. Based on initial price, then, a shopper is unlikely to be able to tell if their plan will deadpool. However, per-person collected premiums in the NAIC filings are much lower for open deadpooling plans than for plans that never close. This reflects differences in pool composition, as enrollees in new risk pools are typically younger and healthier than in more established pools. This difference in composition is also reflected in per-person claims, which are much lower for open deadpooling plans than for plans that never close.

Closed deadpooling plans have annual premiums hundreds of dollars above those of plans that never close in the quoted-price CSG data. The NAIC filings paint a picture of risk pools that have deteriorated, with per-person claims higher than those in plans that never close. Deadpooling plans spend a slightly lower share of their collected premiums on claims both before and after closure. Per-person profits are higher in dollar terms after closure for deadpooling plans.

Table 1. Premiums Rise After Closure for Deadpooling Blocks

	Means by block status			
	Full sample	Never close	Open, DP	Closed, DP
<i>CSG: Quoted rates, 2014–2025</i>				
Annual rate, age 65	\$2,613	\$2,578	\$2,548	\$2,809
Annual rate, age 70	\$3,031	\$2,979	\$2,946	\$3,221
Annual rate, age 75	\$3,404	\$3,289	\$3,363	\$3,729
Annual rate, age 80	\$3,728	\$3,562	\$3,725	\$4,172
Annual rate, age 85	\$3,953	\$3,768	\$4,056	\$4,553
<i>NAIC: Collected premiums, 2006–2025</i>				
Mean per-person profit	\$625	\$543	\$629	\$881
Mean loss ratio	0.80	0.82	0.77	0.78
Mean per-person premium	\$2,976	\$2,852	\$2,504	\$3,507
Mean per-person claims	\$2,351	\$2,309	\$1,875	\$2,626
Mean enrollment	813	1,398	723	338
Plan-block-years	186,561	68,703	11,767	37,714
Unique blocks	17,101	6,658	3,166	4,768
Parent companies	228	141	107	124

Notes: Premiums, claims, and profits are per enrollee and enrollment-weighted. Dollar figures are in January 2025 dollars. The loss ratio is total claims divided by total premiums, and represents the share of premiums collected that are spent on medical claims. Profits are defined as premiums minus claims. “Open, DP” and “Closed, DP” are the pre- and post-closure block-years of deadpooling blocks. “Never close” is the control group of the closure event studies, blocks that never close to new enrollment. Identifying a successor requires five years of post-closure data, so deadpooling blocks are restricted to those closing in or before 2020. Never-closing blocks are retained in full. Unique blocks are counted within each column, so a deadpooling block observed both before and after closure appears in both DP columns, and blocks already closed when first observed appear only in “Closed, DP”. This table pools all regulatory regimes; Appendix A.3 repeats every row separately for open enrollment, annual switching window, and no protections states.

The NAIC filings cover 2006–2025 in 48 states (every state and the District of Columbia, excluding nonstandard Medigap markets in Massachusetts, Minnesota, and Wisconsin). CSG quotes cover all 48 of those states; we use them from 2014 to 2025. Over those years, an own-year CSG quote exists for 82.4% of lives. The CSG row is the annual rate published for one fixed hypothetical non-smoking applicant of the stated age, taken as the equal-weight mean of the male and female quote. We average the CSG rate evenly over ZIP codes within a state. We collapse across companies by weighting with covered lives from the NAIC. A block that neither carries a quote nor appears in the closed-block filing file is considered missing in this table.

5 Empirical Strategy

Our primary estimation strategy looks at the impact of the closure event. Plan closure is our treatment event, and we compare plans that deadpool to a control group of plans that remain open in the market. Because closure is staggered across different years, and we expect treatment effects to grow in each year after closure, we use Callaway and Sant’Anna (2021) as our primary estimator. Alternative specifications for our main results are reported in Appendix C.

5.1 Setup and estimand

We index blocks by i and calendar years by t . Let G_i denote the year in which block i closes to new enrollment, setting $G_i = \infty$ for a block that never closes. We write $Y_{it}(1)$ and $Y_{it}(0)$ for the potential outcomes with and without closure respectively, where Y is per-enrollee premiums, claims, profits, or loss ratios.

A block *deadpools* only if its parent also reopens, so we define the deadpooling indicator:

$$DP_i = \underbrace{\mathbf{1}\{G_i < \infty\}}_{i \text{ ever closes}} \times \underbrace{\mathbf{1}\left\{ \begin{array}{c} i\text{'s parent opens a block in the same state} \times \text{plan-letter} \\ \times \text{market during } [G_i, G_i + 5] \end{array} \right\}}_{i\text{'s parent reopens in the same market within five years}}. \quad (1)$$

For comparison, we report results for nondeadpooling closures, $\{G_i = g, DP_i = 0\}$ in Appendix C.3.

The building block of each estimation is the group-time average treatment effect

$$ATT(g, t) = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid G_i = g], \quad (2)$$

the causal effect on outcomes in calendar year t of plan closure in year g . We use never-closing blocks as the comparison group, under conditional parallel trends: for $t \geq g$,

$$\mathbb{E}[Y_{it}(0) - Y_{i,g-1}(0) \mid G_i = g, X_i] = \mathbb{E}[Y_{it}(0) - Y_{i,g-1}(0) \mid G_i = \infty, X_i], \quad (3)$$

where X_i is a vector of time-invariant controls including plan type and state.

We weight results by enrollment unless otherwise reported to represent the experience of an average enrollee-year. The tables report the overall aggregation of ATTs across years, $\theta = \sum_g \sum_{t \geq g} w_{gt} ATT(g, t)$. The figures report the event-time path $\theta(t - g)$, normalized to the year before closure such that $\theta(-1) = 0$.

We cluster standard errors at the NAIC insurer level throughout, as we find empirical evidence of closure occurring subsidiary-wide (Appendix D.1). We report alternate standard errors clustered at the parent and block levels in Appendix C.4.

5.2 Identifying assumptions for closure

Our design rests on assumptions including parallel trends between treated and control units, that outcomes did not move in anticipation of treatment, and that treating one unit does not change outcomes for comparison units (SUTVA).

Parallel trends. First, we outline our tests of parallel trends. Equation (3) requires that blocks which close would have followed parallel average paths to control blocks in the absence of closure once the covariates X_i are held fixed. This is untestable for $t \geq g$, but we can test whether this is true before closure. We aggregate the pre-treatment cells to event time and test

$$H_0 : \theta(-6) = \theta(-5) = \theta(-4) = \theta(-3) = \theta(-2) = 0 \quad (4)$$

jointly by a Wald test, where $\theta(-1) \equiv 0$ is the base period against which every event-time effect is measured.

The closure event-study figures in Section 6 plot each pre-treatment coefficient against this base period, and Appendix C.1 reports the joint tests.

Consumer anticipation. When blocks close, because existing enrollees are still offered coverage, there is no notification process. From an enrollee’s point of view, deadpooling is close to random. Panel A of Table A4 in Appendix C.2 estimates the annual hazard that an open block deadpools the following year, comparing blocks in the same state, plan letter and year. While plans are still open, their quoted price levels and the growth of these levels, the two characteristics of plans that seniors are most exposed to, do not predict deadpooling. Closure is typically an internal administrative decision about a subsidiary’s whole book, and subsidiaries close blocks across plan letters and states in the same year (Appendix D.1), so closure is a subsidiary-wide decision rather than a response to any single block’s claims.

Interference between units (SUTVA). Two spillovers could contaminate the comparison group. First, enrollees who leave a closed block who are low-risk enough to pass medical underwriting could enroll in other deadpooling or never-closing blocks in the same market. The scope for this type of spillover is limited, as lock-in keeps most enrollees in the closed block (Section 7). Second, closures may move market prices. Section 7.2 finds that markets with more deadpooling have slightly higher minimum quotes. This effect is small, and even if that reflected open blocks increasing their prices in these areas, this would bias our closure effects on premiums toward zero by reducing the gap between treatment and control groups in the most intensely treated areas.

5.3 Legislative analysis

We estimate the effect of states adopting switching legislation using the same Callaway–Sant’Anna estimator as above, with the unit and the timing variables redefined to match the level at which the policy is assigned. Our treatment variable is states passing this legislation, so treatment is assigned at the state $\times$ year level. We use this more aggregated data to look at the total prevalence of outcomes in

the market rather than restricting identification to plans in the market both before and after the regulation is passed. We index state $\times$ year cells by j . Let $s(j)$ be the state of cell j , and let A_s denote the year in which state s adopts a switching rule, with $A_s = \infty$ for a state that never adopts. The group-time average treatment effect is

$$\text{ATT}^{\text{leg}}(a, t) = \mathbb{E} \left[Y_{jt}(1) - Y_{jt}(0) \mid A_{s(j)} = a \right]. \quad (5)$$

This is the causal effect on outcomes in calendar year t of switching law passage in year a . The control group is never-adopting states, with the conditional parallel trends assumption

$$\mathbb{E} \left[Y_{jt}(0) - Y_{j,a-1}(0) \mid A_{s(j)} = a, X_j \right] = \mathbb{E} \left[Y_{jt}(0) - Y_{j,a-1}(0) \mid A_{s(j)} = \infty, X_j \right]. \quad (6)$$

Here, X_j is an optional vector of time-invariant controls including plan type and geographic region. States with switching rights or open-enrollment rules that predate our panel are excluded from the control group. Standard errors are clustered by state, the level at which the law is passed. We observe six treated states, displayed in Figure 2. In light of this limited adoption, the standard errors in Table 6 are wide. We weight results by enrollment unless otherwise reported to represent the experience of an average enrollee-year.

The tables report the overall aggregation across passage years such that $\theta^{\text{leg}} = \sum_a \sum_{t \geq a} w_{at} \text{ATT}^{\text{leg}}(a, t)$. The figures display the event-time paths $\theta^{\text{leg}}(t - a)$, normalized to the year before adoption such that $\theta^{\text{leg}}(-1) = 0$.

5.4 Identifying assumptions for switching laws

Parallel trends. For the state law analysis, the testable implication of (6) is the analogue of (4), over the same five leads,

$$H_0 : \theta^{\text{leg}}(-6) = \theta^{\text{leg}}(-5) = \theta^{\text{leg}}(-4) = \theta^{\text{leg}}(-3) = \theta^{\text{leg}}(-2) = 0. \quad (7)$$

California adopted its switching rule in 2009 and our panel begins in 2006, so it contributes only to the leads at $t - 3$ and $t - 2$. Because only six states adopt a switching rule, conventional clustered inference can over-reject (Conley and Taber 2011; MacKinnon and Webb 2017), so the switching-rule figures in Section 8 also report randomization-inference tests. These reassign the observed adoption years across the eligible states 199 times and recompute the Wald statistic each time.

Law timing and anticipation. Switching laws are passed by state legislatures, and their timing reflects the legislative calendar and the spread of model bills across states more than the condition of any one state's Medigap market. Four of the six adoptions came in 2022, at the beginning of a wave of Medigap switching legislation. We find no evidence that market conditions predict switching law adoption in Table A4, Panel B.

Interference between units (SUTVA). For switching laws, enrollees cannot buy across state lines, so we expect minimal spillovers.

6 How Insurers Profit from Closing Blocks

6.1 What Happens After a Block Deadpools

Table 2 reports Callaway-Sant’Anna (2021) difference-in-differences estimates of the effects of plan closure. Our control group is the set of never-closed plans, and our treatment group is deadpooling plans. Each panel shows the impact of deadpooling closures on a different outcome. Panel A looks at how deadpooling affects premiums, and Panel B looks at medical claims paid out by the insurer. Panels C and D look at firm profit in dollars and the loss ratio respectively, and Panel E looks at enrollment. We report three specifications using different combinations of controls. Column 1 uses no additional covariates, Column 2 adds state fixed effects, and Column 3, our preferred specification, adds state and plan-type fixed effects. The final column reports the enrollment-weighted mean of the outcome before closure for scale. Figure 3 plots the same estimates year by year, normalized to the year before closure. For each outcome below we discuss the table and the matching panel of the figure. Alternate specifications and robustness checks are reported in Appendix C. The same before-and-after-closure comparison under two-way fixed-effects specifications yield similar results (Appendix C.8).

Panel A displays the per-person premium estimates three ways. The top row of this panel uses the NAIC filings. Per-person premiums collected by deadpooling plans rise by \$1,005 after closure in this data. This point estimate is the largest because the NAIC sample reports what companies were actually paid in premiums, not what individual consumers would have been charged. This row includes premium increases due to static pool aging, selection in enrollment over time, and true increases in quoted prices. The aging and selection upward drift is visible in Panel A of Figure 3, where NAIC premiums do not have flat pretrends. The figure also shows that the magnitude of the effect of deadpooling grows each year after closure.

To allow our NAIC estimates to be more comparable with the CSG sample, in the second row of Panel A we restrict the analysis to block-years for which a CSG price also exists. In this sample, deadpooling closures raise collected premiums by \$630. Because the NAIC data begin in 2006, while CSG data begin in 2014, the full NAIC data capture longer deadpooling post periods, which increases the point estimate because treatment effects grow each year after closure. Finally, in the third row of Panel A we report the change in CSG prices associated with a deadpooling closure. This change reflects the difference in unanticipated (not due to predictable aging or inflation) change in prices that the average Medigap enrollee would have faced. Quoted prices for a fixed mix of ages rise by \$435, so repricing of the plan itself accounts for 69% of the increase in premiums collected. That \$435, 15.3% of the pre-closure quote, is the unexpected premium increase the average enrollee faced.

Panel B reports the same Callaway-Sant’Anna specifications, but using average claims paid out by the insurer as the outcome instead of collected premiums. Panel B of Figure 3 shows that deadpooling plans have lower average claims before closure, but substantially higher average claims after closure. This is consistent with low-cost beneficiaries leaving deadpooling plans in larger numbers and concentrating high-risk individuals. In Column 3, average claims rise by \$606 after closure, explaining 60% of the

Table 2. Premiums and Profits Rise After a Deadpooling Closure

	(1)	(2)	(3)	Pre-Closure Means
<i>Panel A: Per-Person Premiums (2025 \$)</i>				
NAIC	\$1,116*** (242)	\$1,049*** (236)	\$1,005*** (235)	\$2,347
NAIC, CSG sample	\$640*** (93)	\$588*** (87)	\$630*** (125)	\$2,203
CSG	\$407*** (70)	\$379*** (65)	\$435*** (78)	\$2,844
<i>Panel B: Per-Person Claims (2025 \$)</i>				
NAIC	\$724*** (139)	\$676*** (146)	\$606*** (139)	\$1,821
NAIC, CSG sample	\$472*** (72)	\$442*** (74)	\$391*** (85)	\$1,723
<i>Panel C: Per-Person Profits (2025 \$)</i>				
NAIC	\$392*** (107)	\$373*** (96)	\$398*** (100)	\$526
NAIC, CSG sample	\$167*** (42)	\$146*** (30)	\$238*** (46)	\$479
<i>Panel D: Loss Ratio (claims / premiums)</i>				
NAIC	-0.035* (0.019)	-0.038** (0.015)	-0.056*** (0.018)	0.795
NAIC, CSG sample	-0.003 (0.014)	-0.003 (0.010)	-0.040*** (0.010)	0.788
<i>Panel E: Enrollment (% of year before closure)</i>				
NAIC	-60.8*** (7.0)	-57.6*** (6.0)	-71.6*** (10.0)	74.5
State controls		X	X	
Plan type controls			X	
N (plan-block-years, deadpooling only)	111,097	111,097	111,097	
Unique blocks	12,478	12,478	12,478	
Subsidiaries	291	291	291	
Parent companies	207	207	207	
N (plan-block-years, CSG sample)	16,928	16,928	16,928	
Unique blocks, CSG sample	2,593	2,593	2,593	
Subsidiaries, CSG sample	118	118	118	
Parent companies, CSG sample	86	86	86	

Notes: Panels report Callaway–Sant’Anna (2021) aggregated ATT of plan closure on each per-person outcome. Each panel reports deadpooling plans only. The parallel estimates for the disjoint set of closing plans whose parent does *not* reopen are in Appendix C.3, alongside these, so the two groups can be compared directly. All monetary units represent annual payments and are adjusted to 2025 dollars. Regressions in Panels A–D are weighted by total lives. Panel E’s outcome is a block’s enrollment as a percent of its enrollment in the year before closure (for never-closing blocks, of their mean enrollment), weighted by that denominator, so each estimate is the change in enrollment in percentage points of year-before-closure enrollment; its pre-closure mean is below 100 because blocks are still growing in the years before they close. *N* counts plan-block-year observations in the estimation sample, not unique plans. Standard errors in parentheses are clustered at the NAIC insurer level. *NAIC* rows use every block-year in the NAIC filings. *NAIC, CSG sample* rows estimate the same outcome only on block-years for which a CSG quoted price also exists (2014 onward), with control block-years limited to blocks quoted in their own right that year; any gap between the two rows reflects the sample, not the measure. The *CSG* row is the CSG quoted-price index: a fixed basket of ages 65–85 weighted by the MCBS age distribution of Medigap holders, carried past a closed block’s last quote by the closed-block rate increases its carrier filed with the state. Counts for the CSG sample are for deadpooling blocks only. All estimates in this table use the doubly-robust inverse-probability-tilting estimator (*drimp*) with a universal base period (*long2*); Appendix C.5 reports the deadpooling premium and profit effects under the alternative doubly-robust estimator (*dripw*). The joint pre-trend test for every panel and column is reported in Appendix C.1, together with what it implies for how each panel should be read. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Joint pre-trend tests (leads $t - 6$ to $t - 2$, $t - 1$ the base period; Appendix C.1): tests reject at the 5% level for premiums ($p = 0.000, 0.001, 0.004$ in Columns 1–3) claims ($p = 0.000, 0.006, 0.025$), and enrollment in Columns 1 and 2 ($p = 0.001, 0.016$); all other pre-trend tests in this table are not significant at the 5% level (profit $p = 0.263, 0.327, 0.188$; loss ratio $p = 0.362, 0.448, 0.260$; enrollment in Column 3, $p = 0.295$).

deadpooling premium increase in the NAIC data. The remainder of the increase in premiums becomes insurer profit.

From Column 3 of Panel C, we see that insurers profit an extra \$398 per person each year after closure in plans that deadpool. Panel C of Figure 3 shows that gap widening in every year after closure as more of the pool is locked in. This increase in profits is not a mechanical scaling up of the firm's original margin: Panel D of Table 2 shows that firms spend 5.6 percentage points less of collected premiums on medical claims after a deadpooling closure.

The annual deadpooling premium rise in column 3 of Table 2, $\Delta P = \$1,005$, can be split into three parts:

$$\Delta P = \underbrace{\Delta C}_{\text{higher claims}} + \underbrace{\Delta C \left(\frac{1}{LR_0} - 1 \right)}_{\text{normal profit margin on the higher claims}} + \underbrace{\left(\Delta P - \frac{\Delta C}{LR_0} \right)}_{\text{markup widening-additional profits}}. \quad (8)$$

Here, ΔC is the rise in claims, and LR_0 is the deadpooling blocks' loss ratio before closure. The portion of increased premiums that go towards paying higher claims for an older and sicker pool is \$606. The second term of Equation (8), \$156, is the margin the insurer would have earned on those higher claims at its pre-closure loss ratio. Without deadpooling, these two terms would still come from premiums, but they would be spread among a wider base of consumers. The third term, \$243, is pricing beyond what the higher claims justify at the old margin. This is the additional transfer from consumers to firms caused by deadpooling.

While firms earn higher margins on each person still enrolled in a deadpooling plan after closure, many people leave deadpooled plans. The enrollment event study in panel E of Figure 3 shows covered lives falling steadily after closure with no pre-trend. That panel expresses a closing block's enrollment as a percent of its enrollment the year before closure, relative to never-closing blocks, using the same Callaway-Sant'Anna design as the closure results. The average enrollee's deadpooling block maintains only 32.5% of its enrollees three years after closure and 15.3% five years after.¹⁵


Why would beneficiaries remain in closed plans? In most states, seniors who have developed medical conditions or are of advanced age are likely to be denied during medical underwriting, leaving them unable to switch their plan without losing access to the comprehensive and flexible coverage offered by Medigap. The concentration of higher-risk beneficiaries in closed plans is consistent with the rise in claims after plan closure. Other beneficiaries may remain enrolled due to inattention.

6.2 Is Deadpooling Profitable?

Section 6.1 established that closing a block raises the profit earned on each enrollee who stays, while the closed subsidiary covers fewer and fewer people. However, total parent profits of deadpooling plans also incorporate gains from the new subsidiary's enrollees. Whether deadpooling pays therefore depends on the parent's combined profit across the old and new books, not on the closed block alone. In this section, we ask whether deadpooling raises that combined profit relative to keeping the original block open. We

15. Each block is weighted by its enrollment in the year before closure, so a block that later empties keeps its weight.

Figure 3. Profit Rises and the Loss Ratio Falls After Closure, as Premiums Outrun Claims and Enrollment Drains



Notes: Panels are per-enrollee premiums (A), per-enrollee claims (B), per-enrollee profit (C), the loss ratio (D), and enrollment as a percent of the year before closure (E), in the order of Table 2. Each panel plots the Callaway-Sant’Anna (2021) ATT of plan closure by event time, aggregated across cohorts and normalized to zero at event time -1 (the year before closure). As in Table 2, the treated blocks are deadpooling blocks and never-treated (never-closing) plans serve as the control group, and the doubly-robust regression adjustment includes state and plan-type (standardized plan letter) controls. 95% confidence intervals use NAIC-clustered standard errors. Dollar outcomes are adjusted to 2025 dollars. Panels A through D are weighted by total lives. Panel E’s outcome is a block’s enrollment as a percent of its enrollment in the year before closure (for never-closing blocks, of their mean enrollment), weighted by that denominator, so a block that later empties keeps its weight. Joint pre-trend tests (leads $t - 6$ to $t - 2$, $t - 1$ the base period): tests reject at the 5% level for premiums, $p = 0.004$; claims, $p = 0.025$; all other pre-trend tests in this figure are not significant at the 5% level (profit, $p = 0.188$; the loss ratio, $p = 0.260$; enrollment, $p = 0.295$).

compare everything a parent earned in a market (state $\times$ plan letter) after its first deadpooling closure with what the same book of business would have earned had the parent not deadpooled.

To create the nondeadpooling counterfactual, we hold the parent's market share at its level in the closure year. We also fix parent's premium level at what it was before closure, letting it grow at the same rate of never-closing blocks in the same state, plan letter, and year to account for increasing medical costs over time. The parent then earns the margin per premium dollar of never-closing companies.

Formally, for a parent company p in market m whose first deadpooling closure is in year t_0 , counterfactual profit $\tilde{\pi}$ in year $t = t_0 + k$ is

$$\tilde{\pi}_{pt} = \underbrace{s_{p0} L_{mt}}_{\text{number of enrollees}} \times \underbrace{P_{p0} \frac{P_{ct}}{P_{c0}}}_{\text{premium per enrollee}} \times \underbrace{\mu_{ct}}_{\text{profit margin}}. \quad (9)$$

Here, s_{p0} is the parent's share of market enrollment in t_0 , and L_{mt} is total market enrollment in year t . Their product is the number of lives the parent would cover if they had maintained their original market share in the year they first deadpooled. P_{p0} is the parent company's premium per enrollee in the year they first deadpooled, P_{c0} is the average premium in never-closing plans from the same year, and P_{ct} is the premium of never-closing plans in year t . $\frac{P_{ct}}{P_{c0}}$, therefore, is the factor by which never-closing premiums have scaled up since year 0. Multiplying the parent's premium level by this ratio yields what the deadpooling parent's average collected premium would be had their premiums grown by the same factor as never-closers. Finally, μ_{ct} is never-closing blocks' profit per premium dollar in year t . The gain from deadpooling is the book's actual profit π_{pt} minus $\tilde{\pi}_{pt}$, summed over $k = 0, \dots, 10$ and discounted to t_0 at 5%.

Deadpooling earns 18.9% more than the counterfactual implies over 10 years, relative to parents that do not deadpool in that market. ¹⁶

6.3 Deadpooling as a Choice: A Natural Experiment in Plan Closure

Deadpooling hinges on a firm's choice to close a block and reprice the pool it retains. What would happen if we compared deadpooling closures to pools that were randomly closed? Using a natural experiment, we can compare the change in premiums and profits after deadpooling closures with the change in those outcomes after mandatory closures. Starting on January 1, 2020, the Medicare Access and CHIP Reauthorization Act (MACRA) barred the sale of two standardized plan types to new Medicare enrollees. ¹⁷ This effectively closed every such block to new enrollees on a single statutory date regardless of the insurer's intent, allowing us to compare deadpooling closures with statutory closures.

16. Each book is followed for ten years after its parent's first deadpooling closure in the market, so only parents whose first deadpooling closure falls through 2015 are included. Summed over the ten years and discounted at 5%, deadpooling books earn 48.7% more than their counterfactual. We apply the same calculation to every parent's book in a market where it never deadpooled, starting it in each year from 2008 to 2015 and weighting the start years to match the deadpooling books. These comparison books earn 25.1% more than their own counterfactual over the same years. The gain from deadpooling is the difference between the two: deadpooling books earn 18.9% more than comparison books, relative to their counterfactuals.

17. Specifically, MACRA banned Plan F and Plan C, to discontinue plans that offered first-dollar coverage and require some nominal cost sharing. At the time of the ban, Plan F was the most popular plan type offered.

In light of the staggered deadpooling closures before 2020, we use Callaway–Sant’Anna event studies to explore this question. For a Plan F or C block the firm deadpooled before 2020 (the same definition of deadpooling as in Table 2), the event time is its filed closure year. For a block MACRA forced shut, the event time is 2020, the single statutory date on which every remaining block of these plan types stopped accepting new 65-year-olds. In both cases, the control group is never-closing, never-deadpooling blocks of other letters, and each ATT is measured relative to event time $t = -1$.

Every force-closed block has event time zero in 2020, whereas the deadpooling blocks are spread across the panel, so the forced arm’s entire post-period coincides with the pandemic and its aftermath. This is a period in which utilization fell and insurer profits rose temporarily in every regulatory regime. Our control group faces the COVID shock as well, so the necessary assumption for this empirical strategy to be valid is that COVID did not differentially affect the newly closed plans compared with other plan types.


Figure 4 displays the effects of deadpooling closure compared to those of statutory closure. Panel A displays profit per enrollee. Profit rises by \$606 for blocks the firm deadpooled, against \$92 for blocks the statute closed. Notably, statutory closures do not exhibit rising profits over time as deadpooling closures do. Panel B displays the effects on quoted prices. This outcome uses a fixed basket of ages, weighted by the MCBS age distribution of Medigap holders, to construct an average quoted premium for Medigap enrollees. Five years after closure, the quoted price is \$348 higher for blocks the firm deadpooled and only \$84 higher for blocks the statute closed, though standard errors are large. Panel C reports the loss ratio; five years after closure, this has changed by -0.103 for blocks the firm deadpooled and by -0.039 for blocks the statute closed. Insurers keep a larger share of each premium dollar after a deadpooling closure than a statutory closure.

In both arms, profit is already rising and the loss ratio already falling in the years before closure, and the joint pre-trend tests reject for both outcomes in both arms. Because the statutory arm shares this pre-closure trend but its profit stops rising after closure, the pre-trend alone does not account for the divergence after closure. The quoted price, which pool composition cannot move, shows no significant pre-trend for deadpooling closures. A captive, aging pool is therefore not sufficient for the post-closure profit increase we observe in deadpooled blocks. This is consistent with deadpooling as a deliberate firm action.

7 Quantifying the Cost to Consumers

In the previous section, we established that deadpooling is strategic and elective firm behavior that raises premiums and profits on enrollees who stay in closed blocks. To quantify the impact on individual enrollees and the total change in consumer surplus resulting from this practice, we must quantify transfers from people who remain in deadpooling plans, switching costs from people who are able to change to other risk pools, and the change in consumer surplus due to loss of coverage among people who are priced out of their plans and do not pass underwriting. The sum of these three lost aspects of

Figure 4. Choice Versus Statute: Profit, Quoted Price, and Loss Ratio



Notes: Plan F/C blocks only. Callaway–Sant’Anna ATT by event time, pooled across regulatory regimes, doubly robust (`drip`) with state controls. Panel A: profit per enrollee (premiums minus claims, 2025 dollars). Panel B: the CSG quoted-price index, a fixed (Laspeyres) basket over ages 65–85 weighted by the MCBS age distribution of Medigap holders, annual 2025 dollars, extended past the end of a closed block’s quotes by the closed-block rate increases its carrier filed; control block-years restricted to those quoted in their own right. Panel C: loss ratio (claims over premiums). Navy dotted: blocks that deadpooled before the 2020 ban (the deadpooling definition of Table 2), on their own closure clock. Red dashed: blocks left open until MACRA force-closed them, on the 2020 clock. Control group: never-closing, never-deadpooling blocks of other plan letters; because no control block is an F or C, plan-letter controls are not identified. Universal base period (long2), $t - 1$ normalized to zero; three pre-closure years are shown. Enrollment-weighted. Standard errors clustered by NAIC company; bars are 95% confidence intervals. Tests that the deadpooling and statutory paths differ add the two arms’ variances and ignore the positive covariance created by their shared control group, so the reported p -values are upper bounds. Joint pre-trend tests (leads $t - 6$ to $t - 2$, $t - 1$ the base period): tests reject at the 5% level for profit, deadpooled, $p = 0.000$; profit, force-closed, $p = 0.000$; the loss ratio, deadpooled, $p = 0.001$; the loss ratio, force-closed, $p = 0.000$; the quoted price, force-closed, $p = 0.002$; the remaining pre-trend test in this figure is not significant at the 5% level (the quoted price, deadpooled, $p = 0.358$).

consumer surplus forms our estimate of the total cost of deadpooling.

7.1 Deadpooling Increases Premiums Unexpectedly

Premiums in most plans increase with age, but deadpooling also raises prices in ways seniors do not anticipate, and these unanticipated increases are large relative to the increases they anticipate for aging. Table 3 separates foreseeable and unforeseeable price increases. For an enrollee aged a in the year before closure, the schedule in that year already states what someone aged $a + 5$ pays, so enrollees have expectations about how their premium will increase as they age. However, they may not expect additional premium increases.

Prices rise over time for enrollees in all Medigap plans. For enrollees in all types of plan, premiums are expected to rise by about \$300 over five years of aging. However, deadpooling plans face additional premium increases of 752, compared to 124 in never-closing plans.

Figure 5 adds controls and shows how the unanticipated part ramps up over time. It displays the closure effect on unanticipated price increases for the average Medigap enrollee, weighted by the MCBS age distribution of Medigap policyholders.¹⁸ This holds the pool constant by construction to isolate the price change an individual would face. Because the ages are held fixed, nothing in this series comes from an enrollee moving along the rate schedule they bought. This is estimated relative to never-closing plans, so increases in the overall price level are also differenced out. This figure then displays the additional premiums that the average enrollee would face due to their plan deciding to deadpool. Five years after

18. This is the series behind the CSG row of Panel A of Table 2. That table averages the effect over all years after closure (\$435 in Column 3).


Table 3. Most Increases in Premiums Are Not Foreseen

	Baseline	Increase over five years		
	premium	Total	Anticipated (age gradient)	Unanticipated (repricing)
<i>Deadpooled blocks</i>				
Aged 65 at closure	\$2,238	\$897	\$305	\$592
Aged 70 at closure	\$2,543	\$1,072	\$376	\$696
Aged 75 at closure	\$2,919	\$1,142	\$347	\$796
Aged 80 at closure	\$3,266	\$1,178	\$313	\$865
<i>Never-closing blocks</i>				
Aged 65	\$2,447	\$776	\$782	\$-6
Aged 70	\$3,229	\$319	\$267	\$52
Aged 75	\$3,496	\$347	\$184	\$163
Aged 80	\$3,680	\$329	\$115	\$214
All deadpooled	\$2,711	\$1,103	\$351	\$752
All never-closing	\$3,229	\$434	\$310	\$124

Notes: A major never-closing insurer gives 65-year-old applicants a large discount that fades as they age, which is why the anticipated increase for never-closing blocks is so much larger at age 65 than at older ages. Annual premiums in January 2025 dollars, lives-weighted by each block's enrollment in the year before closure. For an enrollee aged a in that year, with b the year before closure: the total increase is $p(a+5, b+5) - p(a, b)$; the anticipated part is $p(a+5, b) - p(a, b)$, what the baseline schedule already showed a person five years older paying; and the unanticipated part is $p(a+5, b+5) - p(a+5, b)$, the schedule itself moving. The three are an exact identity. Prices are CSG quoted rates for a non-smoking applicant, carried past the end of a closed block's quotes by the closed-block rate increases the carrier filed with the state; all are the female quote, which is the gender the CSG age profile is built on. Never-closing blocks are measured over the same calendar years, so the deadpooling-specific surprise is the difference between the two unanticipated columns, \$628 per year, not the deadpooled figure itself. A block enters only if CSG quotes it at the enrollee's age in the baseline year *and* at the age they reach five years later, which restricts the sample to closures between 2015 and 2020: 637 of 5,248 deadpooling blocks, but 31.7% of deadpooled enrollment, since the blocks CSG quotes at both endpoints are the larger ones. One never-closing company's Plan A in Oklahoma, whose quote for a 65-year-old rose about sixfold in five years, is dropped as an outlier. The deadpooling-specific surprise here, \$628, differs from the five-year effect of \$762 in Figure 5 because the two use different samples and methods: this table follows enrollees in the 637 blocks CSG quotes at both ages and makes no adjustment for state or plan letter, while the figure is the Callaway-Sant'Anna estimate on the full CSG sample, with state and plan-letter controls.

plan closure, the average Medigap enrollee would be charged \$762 more annually for the same coverage than if their plan had not deadpooled. Part of this premium increase goes towards paying rising claims instead of directly towards insurer profits. To an enrollee, however, that split does not matter; they observe their price increasing to maintain the same benefit, regardless of where the money goes. Therefore, this is our preferred estimate for identifying the effects of deadpooling on the finances of individual enrollees.

Figure 5. The Closure Effect on a Quoted Price, Holding the Pool Fixed



Notes: Callaway–Sant’Anna ATT of plan closure by years since closure, normalized to zero at $t = -1$, with the estimator, covariates, clustering and weighting of Figure 3: never-closing plans as the control group, doubly robust (`drimp`) with state and plan-letter controls, universal base period (1ong2), enrollment-weighted, standard errors clustered by NAIC company; bars are 95% confidence intervals. Outcome: the CSG quoted-price index, a fixed (Laspeyres) basket of ages 65–85 weighted by the MCBS age distribution of Medigap holders, in annual 2025 dollars, carried past the end of a closed block’s quotes by the closed-block rate increases the carrier filed.

7.2 Deadpooling Does Not Make the Market Cheaper

The estimates above price the cost of deadpooling on enrollees in deadpooling plans. However, if closing a block sequesters the sick and lowers prices, deadpooling may redistribute between seniors. This could happen if deadpooling is associated with lower premiums in early years and higher premiums in later years, an intertemporal redistribution within person. It could also be possible that deadpooling reduces the share of sick enrollees in the market and bring down premiums for enrollees in never-closing plans. We investigate this empirically and find no evidence that deadpooling reduces prices in nondeadpooling plans through either mechanism.

Deadpooling is, unsurprisingly, associated with more segmented risk. In markets where deadpooling is more prevalent, per-enrollee claims are pulled further apart across blocks in the same state, plan letter and year. A move from a market with no deadpooling to all deadpooling would raise the between-block coefficient of variation of claims by 25.7%. A more typical difference in deadpooling, from the 25th to the

75th percentile of the deadpooling share of lives, raises it by 3.8%.

However, increased sorting does not make the open blocks cheaper. Comparing CSG quoted prices for open deadpooling blocks and open non-deadpooling blocks, deadpooling blocks are if anything more expensive. Table 4 reports regressions of log prices on an indicator for plan deadpooling. Our preferred specification, in column (4), compares a deadpooling block only with other blocks of the same subsidiary in addition to state, plan type, and year fixed effects. In that comparison, blocks that will deadpool are quoted slightly and insignificantly above their siblings for all consumer ages. The comparisons across companies in columns (1)–(3) give larger positive gaps, and the coefficient is positive in all twelve estimates, so no specification shows deadpooling blocks quoting below the market.

Table 4. Open Deadpooling Blocks Are Not Priced Below the Market


	(1) Unweighted	(2) Lives-weighted	(3) Saturated	(4) Company FE
Age 65	9.9%*** (2.7)	13.6%*** (3.6)	12.5%*** (3.2)	5.7% (3.6)
Age 70	8.0%*** (3.1)	6.3% (5.3)	5.1% (5.5)	2.2% (3.0)
Age 75	7.9%*** (3.0)	7.9%* (4.3)	7.1% (4.5)	2.2% (2.1)
Enrollment-weighted		X	X	X
State, plan-letter, year FE	X	X		X
State × letter × year FE			X	
NAIC company FE				X

Notes: Each cell is the coefficient on an indicator for an open deadpooling block in a regression of the log CSG quoted monthly rate on that indicator, expressed as a percentage difference. Standard errors, in parentheses and also in percentage points, are clustered on the parent (NAIC group). An open deadpooling block-year is a year in which the block is still open to new entrants and the block itself later closes with a successor block opening within five years under the same parent, which restricts treatment to closures through 2020. The comparison group is open blocks that do not meet that definition. Quotes are for non-tobacco applicants and are averaged across ZIP codes within a state. Quotes for men and women are averaged within age. Columns (1) and (2) use additive state, plan-letter and calendar-year fixed effects; column (3) uses a saturated state × plan-letter × year cell, so a block is compared only with blocks quoting the same plan letter in the same state in the same year; column (4) adds NAIC company fixed effects to column (2), so a deadpooling block is compared with the other blocks of the same insurer subsidiary; it is our preferred specification. * $p < .1$, ** $p < .05$, *** $p < .01$. $N = 48,910$ block-years in columns (1) and (2).

To investigate whether areas with more deadpooling have lower prices for enrollees in other plans, we plot the relationship between premiums and the share of the market in deadpooling plans. Among states with no switching protections, a market in which more covered lives sit in deadpooling blocks offers higher price in other plans. A healthy shopper in a deadpooling-heavy market therefore pays slightly more for the cheapest plan available, not less, as seen in Figure 6 (Appendix Figure A11 shows the same plot for the mean quote). This is consistent with a view of deadpooling as a practice that destabilizes the market, perhaps by inducing fewer healthy people to invest in Medigap through reputational damage to the product.

Deadpooling pulls risk apart without passing the proceeds to anyone but the insurer, which is consistent with a setting where demand is price-inelastic. The transfer runs from the sick to the firm, not from the sick to the healthy, so health equity concerns are not offset by a gain elsewhere in the risk distribution.

Figure 6. The Healthy Shopper’s Best Price Is No Lower Where Deadpooling Is Prevalent



Notes: Market = state $\times$ plan letter $\times$ year; outcome is the tenth percentile of quoted rates at 65, in logs. Both axes are residualized on market fixed effects, plan $\times$ year fixed effects, and the number of offers on the shelf. Binned means, twenty equal-size bins, over markets in all states (4,401 markets). Descriptive: deadpooling intensity is not randomly assigned. In this all-state sample the tenth percentile of quoted rates is 10.6% *higher* where deadpooling is prevalent ($p = 0.035$), as is the cheapest offer (10.1%, $p = 0.051$). The estimates in the text restrict to states with no switching protections, where the same comparison gives 6.2 log points for the tenth percentile and 6.6 for the cheapest offer. The equilibrium object a healthy shopper actually faces is therefore no cheaper, and the artifact-free fixed-draw minimum confirms it (5.3% higher, $p = 0.000$).

7.3 Calculating Lost Consumer Surplus: Transfers, Switching Costs, and Loss of Coverage

Section 6.1 separated the premium rise into higher claims, the normal margin on them, and an additional markup. We first aggregate that additional markup, the transfer from people who stay in deadpooled plans to insurers, over the enrollees who pay it. This gives us the total amount of consumer surplus paid in deadpooling. We then calculate how much harm deadpooling does to people who leave their plans. Plans could show decreased enrollment due to mortality, beneficiaries switching to another Medigap insurer, or beneficiaries leaving the Medigap market. After differencing out mortality, total lost consumer surplus is described in the equation below:

$$\begin{aligned} \text{Total lost CS} = & \text{Transfer to insurers from stayers} + \text{Switching cost} \\ & + \text{Lost surplus from those who lose Medigap coverage} \end{aligned} \tag{10}$$

We will estimate each term of this summation. Each term is added over deadpooling closure cohorts g for the five years after closure $s = 1, \dots, 5$. Let N_{gs} be the number of cohort g 's enrollees still enrolled in their deadpooling plan s years after closure and D_{gs} the number who leave in year s after differencing out those leaving the market at the 3.9% annual mortality rate q .

$$\text{Transfer} = \sum_g \sum_{s=1}^5 m_s N_{gs}, \quad (11)$$

$$\text{Switching cost} = \pi \kappa \sum_g \sum_{s=1}^5 D_{gs}, \quad (12)$$

$$\text{Lost surplus} = (1 - \pi) \sum_g \sum_{s=1}^5 D_{gs} v_s \sum_{t=s}^5 (1 - q)^{t-s}, \quad v_s = \frac{1}{2} (e_{s-1} + e_s). \quad (13)$$

First, the transfer to insurers from stayers. The transfer counts only the markup widening of Section 6.1, $m_s = \Delta P_s - \Delta C_s / LR_0$. This is the year- s premium rise beyond what higher claims cost at the pre-closure loss ratio. Part of the higher premiums stayers pay covers the higher claims of a more concentrated pool. If healthy enrollees paid correspondingly less because their pools were marginally less expensive to insure, this would be a transfer from healthy to sick enrollees, but would not be a loss to consumers as a whole. In the sections above, we find no evidence that deadpooling lowers prices for healthy enrollees in plans that do not close; deadpooling plans do not price below the market before they close, and open deadpooling plans quote higher prices to the same applicant. We conservatively measure the total loss to consumer surplus from people who stay in deadpooling plans as the share of premium increases that goes to additional insurer profits. No mortality adjustment is needed for this term of harm, as people who die would be counted among leavers. This gives a total transfer of \$970 million over our period, or \$64.7 million per year of closures. Alternative ways of quantifying this transfer are reported in Appendix E.1.

We next turn to the people who leave deadpooled plans. Five years after closure, if mortality were the only thing removing people from plans, we would expect 82% of the original block to be alive and enrolled.¹⁹ Across the 15 deadpooling closure cohorts, 36.9% of the 2,743,098 enrollees at closure are still enrolled five years later, 17.6% have died, and the remaining 45.5% have left beyond what death explains.

With respect to the Medigap market, these people represent lost total surplus. Excess premiums paid by a senior who stays are transferred to insurers, so no total surplus is destroyed. A senior who gives up coverage because unexpected changes to their premium caused by deadpooling led their Medigap premium to exceed their willingness to pay destroys total surplus in this market. Enrollees lose the risk protection a Medigap provides, and the Medigap insurer does not collect premiums from them. Available data do not allow us to directly observe whether these seniors switch to another Medigap plan, forgo supplemental coverage, switch to Medicare Advantage, or leave the Medicare program altogether.

Leavers are not all harmed in the same way, so we split this group by whether they can obtain another

19. We compute the annual mortality rate for Medigap enrollees to be 3.9%. We calculate this based on the MCBS (Fall 2017–2023, Medigap enrollees, survey-weighted). This gives the age composition of the Medigap population, and the Social Security period life table (Social Security Administration 2024) gives the annual probability of death for ages 65–74, 75–84 and 85 and over. The same rate is used for every mortality adjustment in the paper. A closed block’s pool is sicker and possibly older than the average Medigap pool, so its true mortality may be higher. However, this would also correspond to lower mortality among early leavers, changing overall consumer surplus ambiguously.

Medigap. Using the MCBS underwriting exercise of Appendix B, we estimate π at 26.8%. This is the share of current Medigap enrollees who would pass medical underwriting and could therefore switch to another Medigap, while 73.2% would be rejected in most markets and cannot.²⁰ A switcher pays κ switching cost only one time, when they switch. We set the switching cost at the \$30 median Medigap search cost of Lin and Wildenbeest (2020).²¹ This is not a monetary cost in practice, but includes underwriting paperwork, physician exams, and plan search. We conservatively assume that no enrollees who stay in the deadpooling plan or leave the Medigap market have to pay this switching cost for a failed switch. Switching costs come to \$0.7 million for each year's closures over their first five years, or \$10.0 million for the 2006–2020 closures.

For the 73.2% of enrollees who would not pass medical underwriting but are priced out of their Medigap by large unexpected price increases, we calculate total destroyed surplus. Because it holds the age mix fixed and is measured against plans that never close, e_s is the unexpected change in the price an enrollee faces, net of aging along their own rate schedule and of market-wide price growth. As a closed block's premium ramps up, an enrollee leaves when the price exceeds their willingness to pay.²² We estimate a leaver's willingness to pay as the midpoint between the highest excess premium they paid while enrolled, e_{s-1} , and the one at which they left, e_s ; this midpoint is their annual lost surplus v_s , relative to the premium they would have paid had the plan not closed. A leaver who loses coverage forgoes v_s in every remaining year of the five in which they are still alive, so the last term carries the survival factor $(1 - q)^{t-s}$. Adding up how many people left and how much surplus each of them gave up traces part of a demand curve, and the area under it is the surplus the ramp destroys.

Figure 7 displays the distribution of lost consumer surplus among these enrollees. While most enrollees who are priced out leave before the increase passes \$250, more than a third endure hundreds or thousands of dollars in premium increases before leaving. Counting only departures beyond expected deaths, charging only the premium increase in excess of what open blocks levy, and valuing each leaver at their own revealed surplus, the destroyed surplus is \$43.0 million for each year's closures, or \$645 million for the 2006–2020 closures.²³

Table 5 collects the result. Panel A looks at how the average individual in deadpooling plans was impacted by this practice. For these enrollees, deadpooling led to \$1,445 in lost consumer surplus per enrollee over five years, with harm spread unequally over those still enrolled, those who switched insurers, and those who lost Medigap coverage. Consumers who remain in deadpooled plans for all five years pay \$2,273 in premium increases above aging and whole-market price increases. Switchers pay only a \$30 one-time cost, while people priced out lost \$776 in total foregone surplus.

The aggregate totals in Panel B leave out the higher claims of the remaining pool and the insurer's


20. The distribution is close to bimodal. 73.2% of enrollees are rejected in more than half of the 2022 markets, so this split is insensitive to whether “can switch” requires acceptance everywhere or merely somewhere.

21. Maestas, Schroeder, and Goldman (2009) estimate a larger Medigap search cost of \$72. Switching costs scale one for one with this value, so they remain a small part of the total under either estimate.

22. Exit does respond to price. Across blocks, each additional \$100 of quoted price changes a block's enrollment the following year by -5.20 percent (s.e. 0.76), with block and year fixed effects.

23. Using the minimum possible WTP (the largest price increase absorbed) instead of the midpoint, this would be \$31.5 million for each year's closures. Using the maximum WTP (the price increase the enrollee left at), it would be \$54.4 million.

Figure 7. Most Enrollees Who Are Priced Out Leave Before the Increase Passes \$250



Notes: The distribution of enrollees priced out of a deadpooled block across the size of the premium increase in force when they left, as a share of everyone priced out within the window (77.9% of the closing cohort). Binning is at the *block* level: for each of 747 blocks that closed early enough for seven post-closure years to be observable, we take that block's own premium path net of ordinary open-block growth and its own retention path, so the increase on the horizontal axis is the one that block's enrollees actually faced rather than a fitted average. The window is a calendar requirement, not a survival one—a block whose enrollees have all gone by year four contributes those departures rather than being dropped for having no later rows. Departures are measured net within each block-year; premiums are winsorized at 1/99 and mortality is not netted out. The price on the horizontal axis is the CSG quoted-price index rather than the realized per-enrollee premium, weighted over ages 65–85 by the MCBS age distribution of Medigap holders and carried past the end of a block's quotes by the closed-block rate increases its carrier filed, so a move along an attained-age rate schedule does not count as an increase. The open-block path the increase is measured against is estimated on the same index among never-closing blocks. The harm totals in the text use the average price path of Figure 5 rather than these block-level paths, net of expected deaths, and value the surplus destroyed at \$43.0 million for each year's closures over their first five years.

Table 5. Reduction in Consumer Surplus for the Enrollees of a Deadpooled Block

<i>Panel A: Cost per enrollee, closures CSG quotes (2014–2020), including higher claims</i>				
Cost per person in group (share of cohort)				
Horizon	Repriced	Switching	Loss of coverage	Cohort average
One year	\$151 (80.2%)	\$30 (4.3%)	\$114 (11.7%)	\$136
Five years, cumulative	\$2,273 (41.9%)	\$30 (11.5%)	\$776 (31.3%)	\$1,445
Seven years, cumulative [†]	\$4,268 (29.2%)	\$30 (13.7%)	\$1,339 (37.3%)	\$2,302
<i>Panel B: Total cost, all deadpooling closures in the NAIC filings (\$ millions), markup only</i>				
	Insurer markup	Switching	Loss of coverage	Total
<i>First five years after closure, closures 2006–2020</i>				
Per year of closures	64.7	0.7	43.0	108.3
All 15 closure years	970	10.0	645	1,624
<i>First seven years after closure, closures 2006–2018</i>				
Per year of closures	102.1	0.7	69.5	172.3
All 13 closure years	1,328	9.0	904	2,240

Notes: Cost per enrollee *at closure*—a cohort average over everyone on the block, not the amount any one person bears. Undiscounted, like Panel B; deaths follow Social Security single-age death rates at each block’s own age mix. The first three columns are the three forms the loss takes—excess premium paid while still enrolled, the one-time search cost for those who re-buy Medigap, and the out-of-market loss for those who cannot—each reported as the cost borne by a person *in that group*, with the share of the cohort in the group in parentheses. The remaining 3.7% (one year) and 15.3% (five years) of the cohort have died, so the three shares and the share who died partition the cohort, and the cohort average in the fourth column is total cost divided by everyone enrolled at closure. It is the figure comparable to the per-enrollee estimates elsewhere in the paper. The switching channel is a switching cost and the out-of-market channel is the change in consumer surplus due to loss of coverage; the repriced channel is a transfer out of the stayer’s pocket, part insurer margin and part the higher claims of the sicker pool that remains (Section 7.1). At five years 17% of the total comes from switching costs and loss of coverage. The stayer column is the total excess paid by an enrollee who stays for the whole horizon. Both panels include blocks whose closure year is imputed from their last filing; they cover a small share of lives, 4.8% of lives at closure in Panel A and 4.9% in Panel B. The premium path in Panel A is the Callaway–Sant’Anna estimate of the closure effect on the CSG quoted-price index (Figure 5), by year since closure and applied to every enrollee: \$151 in the first year after closure, rising to \$762 by the fifth. It is measured against never-closing plans and relative to the year before closure. Because CSG quotes begin in 2014, Panel A’s seven-year row can follow only the earlier closures. Panel B instead covers every deadpooling closure from 2006 to 2020 in the NAIC filings (2,743,098 enrollees at closure), each followed for five years, in January 2025 dollars, undiscounted (equations (11)–(13)). Its loss-of-coverage column values leavers on the same Callaway–Sant’Anna price path as Panel A. Its first column is the insurer markup only: the premium rise beyond what the higher claims cost at the pre-closure loss ratio, estimated year by year (equation (11)). It leaves out the higher claims of the remaining pool and the insurer’s normal margin on them, which are included in the per-person figures of Panel A but not in these totals. Each dollar of markup costs enrollees \$1.67. The seven-year rows follow only the closures observed for seven years. [†]Panel A: the 740 blocks that closed in 2014–2018, which CSG quotes for seven years after closure; its shares partition that cohort, of which 19.9% have died. Panel B: the 13 closure years 2006–2018, with markups in years six and seven of \$382 and \$508 per enrollee, estimated with the specification of column 3 of Table 2; over seven years each dollar of markup costs enrollees \$1.69.

normal margin on them, as these impact premium levels but would be spread among consumers regardless of deadpooling. Even so, the additional markup, switching costs, and out-of-market loss together come to \$1.67 per dollar of additional producer surplus.

Overall, 36.9% of enrollees at closure are still enrolled five years later and pay the markup throughout. Some leavers beyond the 3.9% annual mortality pass medical underwriting and switch Medigap insurers (12.2% of enrollees at closure), paying a one-time switching cost of \$30. The other 33.3% of enrollees leave Medigap and lose the plan's value, valued from when they left by revealed preference as above, totalling \$645 million over our sample period. Anyone who leaves before year five pays the markup while enrolled and is then counted as a switcher or leaver. During our period, total lost consumer surplus is \$1.62 billion, or \$108.3 million per year of closures. Each dollar of markup costs seniors \$1.67 in lost consumer surplus. We only count harm for five years after plan closure, so this estimate is conservative.

We report three robustness checks in the appendix: whether the harm should be netted against a lower price paid while the block was open (Appendix E.2), a size-matched control group, the market's largest carrier, which never deadpools (Appendix C.11.1), and valuing leavers on the realized NAIC premium path rather than the quoted price (Appendix E.1).

8 Do Switching Rules Curb Deadpooling?

While most states have no legislation in place to address deadpooling, some states require annual periods when beneficiaries can switch plans. This would allow consumers to leave a closed plan when premiums increase, which would render deadpooling ineffective with perfectly rational and attentive consumers. With imperfect consumers, switching legislation transfers the higher premium burden from beneficiaries who would not pass medical underwriting towards beneficiaries who are inattentive or confused about regulations.


8.1 Comparison Across Regimes

Figure 2 maps the state Medigap regulations in effect during and prior to the period we observe. Open enrollment legislation allows people from outside the Medigap market to purchase a Medigap plan at least annually, while switching legislation allows people already in the Medigap market to switch their policy without medical underwriting at least annually. NY, CT, and ME are the open enrollment states, as they each give Medicare enrollees who do not have a Medigap a chance to join the market at least annually. WA and MO are switching states, as are CA, OR, NV, ID, LA, and IL during the years after they changed their regulations. Using this classification, we can compare market outcomes across regulatory regimes to assess whether states with stronger switching protections exhibit different levels of deadpooling.

Deadpooling is far more common where seniors cannot freely switch plans. Figure 8 shows the share of enrollees in plans that deadpool by state regulatory regime. In open-enrollment states, only about 1.3%

of enrollees are in a deadpooling plan. That share rises to 11.6% in states with an annual switching window and to 19.5% where all switching requires medical underwriting. The ordering is consistent with our lock-in mechanism, as trapping and repricing an aging pool is worth much less where enrollees can leave.²⁴ We next look at changes to outcomes in states that adopted switching rules during the sample before and after the laws passed to achieve a causal estimate of the effects of these rules.

Figure 8. Deadpooling Is Far More Prevalent Where Switching Is Restricted



Notes: Share of enrollees in a plan that deadpools, by state regulatory regime, weighted by covered lives and pooled over the sample through 2020 (the years with a full five-year post-closure lookforward). “Open enroll” states (New York, Connecticut, Maine) guarantee the right to switch Medigap plans at any time; “switching” states (e.g. California, Oregon) provide an annual guaranteed-issue window; “no protections” states leave switching to medical underwriting. Deadpooling covers 1.3% of enrollees in open-enrollment states, 11.6% under annual switching, and 19.5% where there are no protections; dropping states that changed regime over the window leaves this nearly unchanged (1.3%, 11.0%, 19.5%). Because a plan is coded as deadpooling if it *ever* deadpools while the regime is assigned by year, this is a descriptive association rather than a causal estimate. Source: authors’ calculations from NAIC Medicare Supplement experience filings.

8.2 Effects of Switching Rules on Deadpooling and Prices

For our analysis of the effect of switching rules to have a causal interpretation, these regulatory changes would have to have been introduced independent of other market reforms. The estimand and identifying assumptions for this design are stated in Sections 5.3 and 5.4, and Panel B of Table A4 finds no evidence that market conditions predict switching-law adoption. Armour and O’Hanlon (2019) discuss regulatory changes to Medicare Supplement guaranteed issue laws in several states, and they find a number of examples where changes were plausibly exogenous of broader market conditions. Several conversations with agents and policymakers have confirmed that these regulatory changes were not in response to uniquely difficult conditions in the states or years where they were passed.

Table 6 reports the overall ATT for each outcome: per-person premiums (Panel A), claims (Panel B), and profits (Panel C), the loss ratio (Panel D), the share of enrollment in deadpooling plans (Panel E), and the share of enrollment in deadpooling plans that have already closed (Panel F). All five columns pool the six adopting cohorts (CA 2009, OR 2013, and ID/NV/IL/LA 2022) against never-treated controls.

Figure 9 plots the same ATT aggregated by event time and normalized to the year before adoption ($t =$

24. The pattern is nearly unchanged when states that altered their regime over the sample are dropped (1.3%, 11.0%, and 19.5%).

–1). The table reports the average post-adoption effect, while the figure shows its dynamics and pre-trends.

In Figure 9, pre-treatment point estimates are generally small in magnitude and statistically indistinguishable from zero, consistent with the parallel trends assumption. Post-adoption, deadpooling rates decline, while profits, premiums, and claims show no clear effects.

Panel A of Figure 9 looks at the effect of deadpooling on premiums. This is a realized per-enrollee amount, so a null could be a price effect cancelling against a composition effect. This would be the case if quoted prices decreased, but older beneficiaries were able to stay in the market, or if quoted prices increased but more younger beneficiaries joined the market. Figure 10 runs the same design on the CSG quoted-price index, which holds composition fixed, and finds a modest though imprecise decreases in quoted premiums. The effect on claims in Panel B of Figure 9 ambiguous, suggesting that neither route is sufficient to inspire large changes to the overall market.

The clearest effect in Table 6 is on the enrollees the rule is meant to protect. The share of enrollment in deadpooling plans that have already closed falls significantly on both samples, by -2.41 percentage points in the full NAIC data and by -1.40 on the CSG sample (Panel F). This removes between a third and a half of the enrollment in these plans, depending on the exact specification and sample.

With only six treated states, and four of them adopting in 2022, the state-clustered standard errors are wide for most other outcomes. The per-enrollee profit effect is positive but statistically indistinguishable from zero (Panel C), so we find no clear evidence that switching rules change insurer profitability. A premium increase alongside a fall in deadpooling is consistent with switching rules modestly loosening lock-in while admitting some adverse selection into the market. The loss ratio (Panel D), which nets out the enrollment-denominator shift, tells the same story from the cost side.


Taken together, switching rules reduce the number of enrollees in deadpooling blocks, and especially those in closed deadpooling blocks. Premiums, claims, the loss ratio, and profits per enrollee do not move significantly (Panels A–D), and the share of enrollment in deadpooling plans as a whole does not fall significantly (Panel E).

8.3 Do Switching Rules Drive Insurers Out of the Market?

Opponents of switching legislation argue that it makes Medigap unprofitable, thins the set of carriers, and ultimately degrades the market (Claassen 2025; Insights 2025). We check this empirically by seeing whether insurers leave the state after states pass switching rules, and we find no short-run insurer response to the adoption of switching rules, although California and Oregon have fewer carriers in the long run.


We first look at overall averages. Appendix Figure A9 depicts average per-person paid premiums, claims, and profits by regulatory regime. The differences between states under switching regimes and states with no protections are essentially zero for all three of these metrics. This is inconsistent with switching regulation generating high levels of adverse selection. As averages may mask important

Figure 9. Switching Rules Reduce Enrollment in Deadpooling Plans but Leave Premiums, Claims and Profits Unchanged



Notes: Panels follow Table 6. Each panel plots the Callaway–Sant’Anna (Callaway and Sant’Anna 2021) ATT of switching-rule adoption by event time, with never-treated states as the control group, normalized to the year before adoption ($t = -1$, the dotted vertical line). Estimates are weighted by total lives, and dollar outcomes are in annual 2025 dollars. Panel E counts everyone in a plan that deadpools, including those whose block has not yet closed. Panel F counts only enrollees whose deadpooling block has already closed, the population actually exposed to post-closure premium increases, so Panel F is a subset of Panel E rather than an independent estimate. Panels E and F are in percentage points of enrollment, the units of Table 6. The figure shows two estimates. The solid blue lines pool all six adopting cohorts and are plotted through $t = +3$, the last event time every cohort reaches. The dashed red lines re-estimate the same specification on California and Oregon *alone* and are plotted through $t = +7$, so the two can be compared at $t \leq 3$ as well as beyond it. Confidence intervals are offset horizontally to prevent overlap. Joint pre-trend tests on the pooled specification (leads $t - 6$ to $t - 2$, $t - 1$ the base period), state-clustered Wald: tests reject at the 5% level for premiums ($p = 0.000$), claims ($p = 0.003$), profit ($p = 0.037$) and the loss ratio ($p = 0.024$), but not for the deadpooling share ($p = 0.848$). Under randomization inference over the adoption assignment (199 draws), no test is significant at the 5% level (premiums $p = 0.095$, claims $p = 0.250$, profit $p = 0.405$, loss ratio $p = 0.455$, deadpooling share $p = 0.970$).

Figure 10. Switching Rules Leave the Quoted Price Unchanged Too



Notes: The design of Figure 9 run on the CSG quoted-price index—a fixed (Laspeyres) basket over ages 65–85 weighted by the MCBS age distribution of Medigap holders, in annual 2025 dollars, extended past the end of a closed block’s quotes by the closed-block rate increases the carrier filed with the state. Outcomes are collapsed to state $\times$ plan-letter $\times$ year, which is the level at which the switching rule is assigned, and the ATT is normalized to the year before adoption. Unlike the closure design, no block-level provenance screen is applied to the control side: treatment here varies across *states* and every cell on both sides averages open and closed blocks alike, so screening blocks would drop treated and control cells unevenly across states rather than protect the counterfactual. Because the basket holds the age mix constant, a null here cannot be a price increase offsetting a healthier pool. Unlike Figure 9, there is no separate California–Oregon line: the CSG quotes begin after Oregon’s 2013 adoption, so neither state has a pre-adoption quote to normalize against, and the pooled estimate here rests on the four 2022 adopters. Joint pre-trend tests on the pooled specification (leads $t - 6$ to $t - 2$, $t - 1$ the base period), state-clustered Wald: no pre-trend test in this figure is significant at the 5% level (the quoted-price index, $p = 0.496$). Under randomization inference over the adoption assignment (199 draws), no test is significant at the 5% level (the quoted-price index, $p = 0.940$).

Table 6. Switching Rules Reduce Enrollment in Closed Deadpooling Plans

	(1)	(2)	(3)	(4)	(5)	Pre-adoption mean
All treated states pooled						
Panel A: Per-Person Premiums (2025 dollars)						
NAIC	57.27 (105.45)	49.23 (102.49)	25.64 (106.50)	49.52 (102.76)	85.56 (87.03)	3,097
NAIC, CA and OR only	89.52 (163.32)	82.85 (158.60)	62.88 (165.24)	83.31 (158.98)	180.48 (110.91)	2,845
NAIC, CSG sample	21.50 (45.81)	22.47 (35.70)	-6.14 (32.04)	22.67 (35.31)	-3.07 (29.38)	2,984
CSG	-145.20 (124.35)	-142.89 (126.03)	-165.08 (127.58)	-143.05 (125.93)	-164.87 (128.45)	3,251
Panel B: Per-Person Claims (2025 dollars)						
NAIC	16.27 (41.36)	10.97 (42.31)	-3.08 (48.51)	10.63 (42.04)	35.62 (58.93)	2,358
NAIC, CA and OR only	-47.59** (23.68)	-55.30** (23.26)	-72.47** (34.34)	-55.51** (23.22)	-27.73 (65.09)	2,239
NAIC, CSG sample	231.28*** (83.24)	233.97*** (84.20)	203.32** (84.06)	233.74*** (83.32)	204.91** (87.28)	2,303
Panel C: Per-Person Profits (2025 dollars)						
NAIC	41.00 (112.89)	38.26 (114.09)	28.72 (126.94)	38.89 (113.94)	49.94 (139.39)	739
NAIC, CA and OR only	137.12 (156.26)	138.15 (157.09)	135.34 (175.89)	138.82 (156.93)	208.21 (170.77)	606
NAIC, CSG sample	-209.78* (113.30)	-211.51** (99.70)	-209.46** (97.18)	-211.06** (98.17)	-207.97** (94.63)	680
Panel D: Loss Ratio (claims / premiums)						
NAIC	-0.021 (0.034)	-0.021 (0.034)	-0.026 (0.041)	-0.021 (0.035)	-0.043 (0.040)	0.768
NAIC, CA and OR only	-0.049 (0.047)	-0.050 (0.048)	-0.061 (0.056)	-0.050 (0.048)	-0.103*** (0.032)	0.805
NAIC, CSG sample	0.044*** (0.011)	0.045*** (0.008)	0.041*** (0.010)	0.045*** (0.008)	0.041*** (0.010)	0.776
Panel E: Share in Deadpooling Plans (percentage points)						
NAIC	-2.89 (1.80)	-2.97 (1.84)	-4.03** (2.01)	-2.99 (1.85)	1.61 (8.35)	16.38
NAIC, CA and OR only	-4.16 (2.62)	-4.25 (2.68)	-5.92** (2.73)	-4.29 (2.69)	3.34 (14.87)	24.13
NAIC, CSG sample	-0.81** (0.33)	-0.83** (0.33)	-0.76* (0.45)	-0.82** (0.33)	-0.77* (0.44)	13.29
Panel F: Share in Closed Deadpooling Plans (percentage points)						
NAIC	-2.37*** (0.58)	-2.41*** (0.57)	-2.45*** (0.47)	-2.47*** (0.61)	0.05 (3.00)	6.87
NAIC, CA and OR only	-3.04*** (0.68)	-3.10*** (0.64)	-3.08*** (0.37)	-3.19*** (0.69)	1.13 (5.32)	4.56
NAIC, CSG sample	-1.40*** (0.31)	-1.40*** (0.30)	-1.35*** (0.31)	-1.40*** (0.30)	-1.36*** (0.33)	5.28
Never-treated control	X	X	X	X	X	
Plan-type adj.		X	X	X	X	
Census region adj.			X		X	
Pre-period size				X	X	
	NAIC	CA and OR only	CSG sample			
N (state-plan-yr cells)	8,072	7,320	4,833			
Block-years	200,812	180,614	80,286			
Parent companies	232	224	136			
Subsidiaries	315	303	201			
Life-years	154,736,712	139,600,958	98,513,819			

Notes: Callaway–Sant’Anna (2021) overall ATT of switching legislation, with never-adopting states as the control group, estimated by inverse probability tilting with a universal base period. Each panel reports the outcome on two samples, plus California and Oregon, to match Figure 9. “NAIC” rows use every block-year in the NAIC filings. “NAIC, CSG sample” rows use the same NAIC outcome restricted to the block-years for which CSG also records a quoted price, which are concentrated in 2014 onward and in the larger, actively marketed blocks. “NAIC, CA and OR only” rows re-estimate the NAIC row treating only California (2009) and Oregon (2013). The four 2022 adopters are dropped and the never-treated controls are unchanged. The “CSG” row of Panel A replaces the realized premium with the CSG quoted-price index, a fixed basket over ages 65–85 weighted by the MCBS age distribution of Medigap holders and extended past the end of a closed block’s quotes by the closed-block rate increases the carrier filed. Because the basket cannot move with the pool, this is a price measure free of composition effects. Treatment is assigned at the state level. Panels A–C report annual per-person payments adjusted to 2025 dollars; Panel D reports the loss ratio (total claims over total premiums, winsorized 1/99); Panel E reports the enrollment share in deadpooling plans, in percentage points; and Panel F reports the enrollment share in deadpooling plans that have already closed, in percentage points. All six panels share an identical specification and differ only in the outcome. The panel is collapsed to state $\times$ plan-type $\times$ year with enrollment-weighted cell means, matching the state-level treatment. Standard errors clustered by state. Joint pre-trend tests on the pooled specification (leads $t - 6$ to $t - 2$, $t - 1$ the base period), state-clustered Wald: tests reject at the 5% level for premiums ($p = 0.000$), claims ($p = 0.003$), profit ($p = 0.037$) and the loss ratio ($p = 0.024$). All other pre-trend tests are not significant at the 5% level (deadpooling share $p = 0.848$, CSG quoted-price index $p = 0.496$, premiums on the CSG sample $p = 0.746$). Under randomization inference over the adoption assignment (199 draws) no test is significant at the 5% level (premiums $p = 0.095$, claims $p = 0.250$, profit $p = 0.405$, loss ratio $p = 0.455$).

distributional changes, this plot alone cannot distinguish between offsetting reallocations within the market and a true absence of regulatory effects. However, the higher levels of paid premiums and paid claims in the open enrollment states is consistent with adverse selection into the Medigap market under these regulations.

To explore insurer behavior, we first look at subsidiary withdrawal. As opposed to closure, where a company does not accept new enrollees but continues to offer coverage to existing enrollees, withdrawal is a true market exit. The NAIC filings carry a withdrawal date separate from the closure date, which we use to build a plan-level withdrawal indicator. We then run the same Callaway–Sant’Anna switching-adoption design used above. The lives-weighted withdrawal rate does not move: the overall ATT is essentially zero, with flat pre-adoption trends and no significant effect at any horizon. Counting active subsidiaries and parents with an open block per state-year and running the same design, we find no decline over the window in which all six adopting cohorts contribute. Over the longer horizon available only for California and Oregon, the number of carriers does fall, with the estimated change is -6.1 parents, about 20% of their pre-adoption mean, and -8.2 subsidiaries. Withdrawal and closure rates do not move in either sample (Table 7). A purely cross-sectional comparison across regimes does show fewer carriers under more protective rules. Open-enrollment states average 14 parents with an open block against 31 in switching states and 34 in unprotected states. However, within a state, switching adoption does not shrink the carrier count in the first years after adoption. California and Oregon do have fewer carriers a decade later, but with only two states we cannot separate that decline from other changes in those markets over the same period.

Table 7. Switching Rules and Insurer Exit

	Effect of adoption	Pre-adoption mean
<i>Panel A: Withdrawal rate (% of enrollees)</i>		
All adopters, years 0 to +3	0.066 (0.070)	0.159
California and Oregon, all post-adoption years	-0.005 (0.095)	0.086
<i>Panel B: Closure rate (% of enrollees)</i>		
All adopters, years 0 to +3	1.084 (1.292)	4.773
California and Oregon, all post-adoption years	1.470 (1.170)	3.474
<i>Panel C: Subsidiaries with an open block (count)</i>		
All adopters, years 0 to +3	-0.10 (1.05)	47.58
California and Oregon, all post-adoption years	-8.16*** (2.51)	39.80
<i>Panel D: Parent companies with an open block (count)</i>		
All adopters, years 0 to +3	-0.05 (0.70)	36.11
California and Oregon, all post-adoption years	-6.07*** (0.66)	30.40

Notes: Callaway–Sant’Anna estimates of switching-rule adoption, with never-adopting states as the control group, a universal base period, and states with pre-existing switching or open-enrollment rules excluded, as in Table 6. “All adopters” rows pool the six adopting states and average event years 0 to +3, the window every cohort reaches. “California and Oregon” rows treat only those two states and report the overall ATT across all post-adoption years, as in the dashed series of Figure 9. Panels A and B are the shares of enrollees in plans that withdraw from the market or close to new enrollment that year, computed on state $\times$ plan-type $\times$ year cells with enrollment-weighted means. Withdrawal is a true market exit, taken from the withdrawal date in the NAIC filings. Panels C and D count the subsidiaries and parent companies with at least one open block in each state and year. “Pre-adoption mean” is the mean of the outcome in the treated states before adoption. Standard errors clustered by state in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Taken together, switching rules loosen lock-in without an immediate exit of carriers. The longer-run decline in California and Oregon is the one sign of the thinning market that opponents predict, but the evidence on insurer response is mild and not definitive.

8.4 Who Uses a Switching Rule?

To explore the fall in the closed-block deadpooling share in Panel E and Panel F of Table 6, we look for evidence of enrollees who are likely to shop using this switching right. We ask whether enrollees in closed blocks respond more to unexpected premium increases once a switching rule is in force. We measure the unexpected increase as the year-over-year change in the block's CSG quoted-price index, which prices a fixed basket of ages. Because the age mix is held constant, the increase an enrollee can anticipate from moving along the block's own age schedule is excluded, and what remains is the schedule itself moving. We regress the percent of a closed block's enrollees lost in a year on indicators for the size of the premium increase and interactions with an indicator that the state's switching rule is in force, with state $\times$ plan-type $\times$ year and years-since-closure fixed effects:


$$A_{it} = \sum_k \left[\underbrace{\beta_k \mathbb{1}\{x_{it} \in k\}}_{\text{response to an increase in bin } k} + \underbrace{\gamma_k \mathbb{1}\{x_{it} \in k\} R_{s(i)t}}_{\text{added response under a switching rule}} \right] + \underbrace{\lambda_{s(i)\ell(i)t}}_{\text{state } \times \text{ plan type } \times \text{ year}} + \underbrace{\delta_{t-G_i}}_{\text{years since closure}} + \varepsilon_{it}. \quad (14)$$

Here, $A_{it} = 100(1 - L_{it}/L_{i,t-1})$ is the percent of closed deadpooled block i 's enrollees lost in year t . x_{it} is the year-over-year change in the log of its CSG fixed-age price index, and the bins k are increases of 0–5%, 5–10%, and 10% or more (where years with no increase or a decrease are the omitted group). Increases of these sizes are common: 22.7% of enrollees in closed deadpooled block-years see a rise of 10% or more in the CSG fixed-age price. R_{st} indicates that state s 's switching rule is in force, and δ_{t-G_i} are years-since-closure fixed effects to account for differential mortality by plan age. Observations are weighted by the prior year's enrollment. Because the state $\times$ standardized plan type $\times$ year effects are included, the comparison is between closed blocks in the same market and year that faced increases of different sizes. β_k is the additional share of enrollees lost after an increase in bin k without a switching rule, and $\beta_k + \gamma_k$ is the additional share of enrollees lost after an increase in bin k with a switching rule.

Figure 11 reports the result. Without a switching rule, enrollees leave closed deadpooled blocks only modestly more after larger increases. Relative to a year with no increase or a decrease, the block loses an additional 1.9, 3.4, and 4.6 percentage points of its enrollees after increases of 0–5%, 5–10%, and 10% or more respectively. Where a switching rule is in force, enrollees are more responsive to the same increases, with departure rates of 7.7, 9.1, and 5.6 percentage points. For increases below 10%, the difference between the rule and no-rule responses is significant at the 1% level under state-clustered standard errors; for increases of 10% or more it is not ($p = 0.401$). Appendix E.4 adds block fixed effects, so that each block is compared only with itself in other years. This flattens the no-rule response, and the rule difference becomes significant in all three bins.

While enrollees are more responsive to changes in premiums when they have access to switching rules,

Figure 11. Under a Switching Rule, Enrollees Leave Closed Deadpooled Blocks When Premiums Rise



Notes: Closed deadpooled blocks, individual policies, excluding states with pre-existing switching or open-enrollment rules. Bars are coefficients from equation (14): grey bars plot β_k and blue bars plot $\beta_k + \gamma_k$, each the additional percent of enrollees lost relative to a year with no increase or a decrease. The state $\times$ plan-letter $\times$ year effects absorb the main effect of R_{st} . Observations are weighted by the prior year's enrollment; 95% intervals use state-clustered standard errors, which rest on six treated states. $N = 6,841$ block-years. Appendix E.4 adds block fixed effects and reports the linear version with randomization inference, the specification identified only from the 2022 adopters, the unweighted estimates, and the estimate without Carrier A.

these heightened responses are still small next to what a fully attentive enrollee would do. In switching-rule states, every enrollee hit by such an increase had a cheaper open plan of the same letter at their age, on average \$1,413 a year, or 40% cheaper.

9 Model and Policy Simulation

Section 8 shows that switching rules help the enrollees who use them, but we find little evidence that they change how firms behave. The policies and market changes we would most like to evaluate, such as a direct ban on closing blocks, a more attentive consumer population, or automatic switching, have never been tried, so there is no natural experiment to study. We therefore build a simulation model of the Medigap market, calibrate it to observed moments, and compare these policies holding the same consumers and firms fixed. Section 9.1 describes the model, and Sections 9.2 and 9.3 report the counterfactuals.

9.1 Model Setup

We build a simulation of the Medigap market to allow us to comment on mechanisms and other policy solutions. We calibrate the model to moments of the observed market and then simulate counterfactual policy regimes from the calibrated steady state.

The market consists of a set of F parent insurance companies, each operating across S states and offering L standardized plan types. Claims and premium levels are calibrated so that the simulated market in aggregate matches the Medigap market during our period. Each company–state–plan type defines a

group of enrollees sharing a risk pool. The market begins with $F \times S \times L = 555$ blocks, and blocks open and close over time as insurers enter, exit, and deadpool. Appendix G.8 gives the values we use and shows that our results are insensitive to raising the number of companies and states.

The model has two types of agents. Consumers enter the market at 65 and are described by five characteristics. First, they have a health state of healthy, moderate, or sick, which can only worsen over time. Second, consumers have annual claims, which are drawn from their health state and age. Third is a willingness to pay for coverage, which is drawn from a wide distribution centered on the certainty equivalent of expected future claims for someone in their health state.²⁵ This lets consumers in the same health state differ in how much they value coverage. Fourth is a price-sensitivity index that governs how often consumers shop.²⁶ Finally, consumers have a mortality risk that rises with age and illness.

These characteristics are drawn from data. Relative claims by health state come from the MCBS Cost Supplement, the age profiles of claims and mortality from the MCBS and the SSA period life table, the shape of the price-sensitivity distribution from our UAS module, which shows that most enrollees rarely shop, with its scale set so that the average shopping rate matches survey evidence on how many beneficiaries compare coverage options each year (Ochieng et al. 2024). Underwriting pass rates by health state come from the MCBS (Appendices G.2, G.3, and B).

The other type of agent is firms. Firms are parent companies, each owning one block in every state and plan type. Every block sets its premium as a markup over its own recent claims. A parent is either assigned to be a deadpooling type or not. Deadpooling parents' blocks follow the standard markup schedule until they close. Closure is based on block age, and the age at which a block closes is drawn from the closure-age distribution in the NAIC panel. We assign the deadpooling type and the closure timing rather than letting firms choose them because closures in the data are coordinated, subsidiary-wide decisions rather than responses to any one block's pool (Appendix D.1). Modeling closure as an optimization would require the costs that keep most parents from deadpooling, such as regulatory scrutiny or reputational damage, none of which we observe. Taking the closure timing from the data lets the model reproduce deadpooling as it occurs without inventing those costs. The share of deadpooling parents and the steepness of the markup schedule are calibrated in Appendix G.9, and Appendix G.8 shows that the results are insensitive to the number of companies and states.


The model period is one year, and Figure 12 sets out the order of events. Insurers move first by posting premiums based on a markup of their prior year's claims. In the initial market, every block opens at a common premium. After premiums are set, enrollees age, switch health states, and realize mortality based on their age and health state. Enrollees switching health states redraw their claims and willingness to pay from their new distribution. In the same step, new 65-year-olds enter the market.

Next, consumers, including that year's entrants, act on the posted premiums: they stay in their current plan, switch to a cheaper open block if they pass underwriting, or exit the market. Even a consumer who

25. This is based on expected average annual claims over a twenty-year forward simulation, plus a risk premium reflecting aversion to the variance in those claims.

26. The index is drawn as $\lambda_i \sim \text{Beta}(0.60, 2.40)$, placing most mass near zero. We justify this choice using results from a study we ran with the Understanding America Study: 68% of UAS enrollees consulted only one agent at enrollment, and 73% have never switched, suggesting that many people do not shop aggressively in this market. This is consistent with Ochieng et al. (2024).

Figure 12. The Order of Events Within a Model Year



Notes: The figure reads top to bottom as one model year, in the order the model executes it. Insurers move first by posting premiums (dashed boxes): in the initial market every block opens at a common initial premium, and in later years consumers face the premiums set at the end of the previous year. Enrollees then age, their health may worsen ($H \rightarrow M \rightarrow S$), and mortality is realized; new 65-year-olds then enter. Consumers then stay, switch to a cheaper open block if they pass underwriting, or exit. Claims are then realized, and insurers reset premiums (dashed box): each block reprices as a markup over the claims it has just paid, and a deadpooling parent may close a block and open a successor. A closure therefore first reaches enrollees' premiums the following year, and because the enrollees best able to leave a closed block are the healthy, the pool that remains deteriorates. Consumers out of the market may re-enter if a premium falls below their willingness to pay and they pass underwriting. Underwriting pass rates (ρ_H, ρ_M, ρ_S) fall a further point per year of age above 65, so sick and older enrollees are effectively locked in.

shops switches only if the annual savings are large enough: the threshold is \$75 divided by their price sensitivity, or \$375 a year for a consumer of average price sensitivity.²⁷ Claims are then realized, and insurers set the next year's premiums based on the claims of their enrollees. A deadpooling parent may also close a block and open a successor.

A consumer who is out of the market faces a different decision: they may re-enter if a premium falls below their willingness to pay, but re-entry is also subject to underwriting, so a consumer who has fallen ill since leaving may be unable to return. Sick consumers are effectively locked into the plan in which they first enroll, by both high WTP and low underwriting pass rates.

Each block sets a community-rated premium as a markup over its own most recent claims experience:

$$P_b^{t+1} = m_b \cdot \bar{c}_b^t, \quad (15)$$

where $\bar{c}_b^t$ is block b 's average claims per enrollee in the preceding year. Blocks of non-deadpooling parents all use a common markup, pinned from the observed premium-to-claims ratio. For deadpooling parents, new blocks are priced slightly above the common markup in their first few years, and once a block closes its markup ramps up over three years and then holds at a higher level, with the size of the ramp calibrated to the data. A new block opens at a fixed initial premium and is then repriced on its own claims. Blocks under deadpooling parents close at an age drawn from the distribution of closure ages in the NAIC panel and applied only to the 11% of parents designated deadpooling types. Once closed, healthy consumers exit to open blocks while sick consumers remain trapped, deteriorating the risk pool, raising premiums, and deepening the adverse selection spiral. Additional model details are in Appendix G.

The model is calibrated by grid search to ten moments of the Medigap market: the average premium and average claims of open and closed deadpooling blocks, the share of enrollees sitting in a closed deadpooling block, the four-point post-closure retention curve, and the annual rate at which enrollees leave Medigap. Four parameters are free (the share of parents that deadpool, the scale of the markup ramp for deadpooling parents, and the responsiveness and convexity of exit in the price). The open-block markup is pinned to match the data. Table A18 in Appendix G.9 reports the fit, and Appendix G.9 gives the targets and the three primitives we re-derive.

9.2 Counterfactual regimes and aggregate effects

To complement the empirical evidence, we take the calibrated model as described above and change only the rules of the market, leaving every primitive, parameter and calibration target fixed. Each counterfactual below is therefore the same market and the same consumers under a different regulatory

27. An enrolled consumer shops with a probability that rises with their price sensitivity. A consumer who shops compares their premium with the cheapest open block of the same plan type in their state, and switches if the savings exceed a threshold that falls with their price sensitivity and they pass medical underwriting. A consumer exits the market with a probability that rises with their price sensitivity and with how far their premium exceeds their willingness to pay. New 65-year-olds buy without underwriting. An entrant who shops buys the cheapest open block if its premium is below their willingness to pay. An entrant who does not shop is assigned to an open block with probability inversely proportional to its premium, and is redrawn among the remaining blocks until they reach one priced below their willingness to pay. An entrant for whom every open block costs more than their willingness to pay stays out of the market.

regime, and any difference between regimes is attributable to the rule change. We report steady-state effects. The baseline calibrates to match empirical moments from states without guaranteed-issue or birthday-rule protections on the 2025 NAIC panel. We compare a demand-side switching policy, a birthday rule (underwriting waived for plan switches once per year); a supply-side ban on deadpooling (a direct prohibition on closing blocks to new enrollment); and consumer-attention counterfactuals in which every consumer's price sensitivity is scaled up (savvy consumers), alone and combined with the birthday rule. We do not model continuous guaranteed issue, under which anyone could buy Medigap at any time without underwriting, because that counterfactual requires additional assumptions about the Medicare Advantage market for which our data are not well suited. We also simulate automatic switching, where each year the regulator reassigns an enrollee to the cheapest open block of the same plan letter in their state, with underwriting waived, whenever that block's premium is at least 10% below what they currently pay. There is precedent for automatic switching in Medicare. In Part D's Low-Income Subsidy program, enrollees whom Medicare assigned to a plan are automatically reassigned to another plan if their plan's premium rises above the subsidy benchmark (Decarolis 2015).


At baseline, 10.8% of enrollees are in a closed deadpooled block. An outright ban on closures empties those blocks almost entirely. A birthday rule does considerably less, cutting the trapped share to 9.5%. A switching right is only as effective as the enrollee's ability to use it, and at baseline attention most trapped enrollees do not shop even once the underwriting barrier is lifted. Aggregate premiums and profits move very little under these policies: the average premium changes by $-\$5$ under the birthday rule and $-\$13$ under the ban, against a baseline of $\$2,862$, while per-enrollee profit stays within $\$20$ of its $\$571$ baseline. Market-wide averages are thus nearly unchanged even where the small proportion of affected high-risk enrollees gain sharply. The consumer-attention counterfactuals move the aggregates most. Raising price sensitivity alone lowers the average premium by $-\$178$ and profit by $-\$43$.

9.3 Where the Harm Goes and Whom It Hits

We evaluate each reform on two measures. The first, the share of trapped-enrollee harm, asks how effective a reform is at eliminating deadpooling, without considering other effects it may have on the market. For each person in a closed deadpooling block, we measure harm as the difference amount people overpay, or the amount of surplus lost from leaving the market, that can be attributed to deadpooling. Because the benchmark moves with each block's own claims, an older or sicker pool does not register as harm. Figure 13 reports this harm as a percent of today's level. The measure is zero by construction when no block deadpools, so a ban on deadpooling removes it entirely. A birthday rule leaves 74% of harm, and attentive consumers 52%, while attention combined with a birthday rule leaves 4% and automatic switching 3%. Appendix G.10 gives the details.

Most harm from deadpooling comes from increased premiums paid by stayers in this model. However, in the attentive consumers counterfactual, a substantial share of harm comes from seniors being priced out. Attentive consumers shop more often, but underwriting is still binding for sick enrollees, so their concentration into deadpooling blocks happens much faster than with inattentive consumers.

Figure 13. Share of Trapped-Enrollee Harm Remaining Under Each Reform




Notes: The share of trapped-enrollee harm: harm caused by deadpooling only. Figure 14 instead reports overall consumer welfare, lost consumer surplus from all sources. Calibrated model, steady state (years 60–79), four seeds, the same run and regimes as Figure 14. Harm from deadpooling alone as a percent of today's level, split into stayers' premium above the open-market markup, the lost surplus of people who left a closed deadpooling block but would stay at that markup, and switching costs out of closed deadpooling blocks. Zero by construction when no block deadpools.

This first measure shows how effective each reform is at getting rid of deadpooling, but it considers only enrollees in closed deadpooling blocks. A reform could remove that harm and still have unintended consequences elsewhere in the market, including for healthy shoppers, enrollees in open blocks, and those priced out of Medigap. Our second measure, overall consumer welfare, measures total lost consumer surplus, scoring every living beneficiary in the model against a common baseline, the cheapest open premium r in their market under the initial rules. Harm is the surplus a person would have at that price, $\max(\text{WTP} - r, 0)$, minus the surplus they have. This splits exactly into four sources: people priced out of a plan they value above r ; enrolled people paying more than the cheapest open plan who shopped and were rejected by underwriting (locked in); enrolled people paying more than the cheapest open plan who did not shop (inattention); and the change in open premiums themselves.

Under the baseline, 14% of harm is lock-in, 57% is inattention, and 9% is overpayment by people who shop without moving. Switching rights remove lock-in, but total harm falls by only 5%. The enrollees that underwriting used to hold in place now switch, but shoppers choose the plan quoting the lowest rate, and the lowest premiums are charged by young blocks whose healthier pools keep their premiums below their quoted rates. Those who switch therefore still pay well above the cheapest open premium, so overpayment by people who shop rises from 9% to 19% of today's harm, and inattention is essentially unchanged (56%). A ban on deadpooling does not reduce total harm, as lock-in falls but higher premiums in open blocks cancel out these gains. Automatic switching cuts harm by 53% with no change in attention and lowers average prices by keeping more low-cost enrollees in the market. Increasing attention alone cuts it by 11%, but attention combined with a birthday rule cuts it by 83%. This solution prevents people from being priced out while keeping healthy enrollees in the market, decreasing the overall price level. Appendix G.10 reports harm by the split of overpayment between lock-in and

inattention by each enrollee’s probability of passing underwriting rather than by what they did. In Appendix G.10, we ask which of these policies the model’s population would vote for, to see the disbursal of harm and benefit across enrollees in addition to overall averages. Table A20 in the same appendix reports harm per affected person rather than per beneficiary.

Figure 14. Overall Consumer Welfare: Lost Consumer Surplus per Beneficiary Under Each Reform, and Where It Comes From



Notes: Overall consumer welfare: total lost consumer surplus from all sources, not only deadpooling (Figure 13 reports the share of trapped-enrollee harm), per living Medicare beneficiary per year in the calibrated model, as a percent of today’s total, steady state (years 60–79, four seeds). Each person is scored against the baseline’s cheapest open premium r in their state, plan letter and year; $\text{harm} = \max(WTP - r, 0) - CS$. Each segment is one source’s total divided by all living beneficiaries, so the segments sum to the total (diamond, with the percent change from today). *Locked in*, *Inattentive* and *Shopped, still above cheapest*: premium paid above the cheapest open plan by people who that year shopped and were rejected by underwriting, did not shop, or shopped but did not move to the cheapest open premium (their savings fell below their switching threshold, or the plan quoting the lowest rate does not charge the lowest premium). *Priced out*: out of Medigap with $WTP > r$, by current health. *Prices vs. today*: change in the cheapest open premium, negative when open plans are cheaper than today.

10 Conclusion

The Medicare Supplement market insures almost 14 million beneficiaries in plans that cost thousands of dollars a year out of pocket. Because Medigap is a comprehensive insurance product sold to a population whose health deteriorates relatively quickly, most states shield it from adverse selection through lock-in. In practice, this means that after an initial enrollment window, insurers may underwrite applicants, so a beneficiary who falls ill cannot readily move. While limiting adverse selection, lock-in also creates ex-post market power for the firm a beneficiary first enrolls with. In this paper, we study the consequences of that market power.

To circumvent existing federal regulation, parent firms close a block to new members, let the surviving pool age and sicken, and raise its price against enrollees who can no longer pass underwriting elsewhere. They then open a fresh block through a different subsidiary to gather the healthy applicants

the closed block can no longer attract. We show that closure raises premiums well beyond what claims justify, that the increase is larger in a natural experiment that allows us to compare deadpooling and statutory plan closures, and that closures arrive as coordinated, subsidiary-wide events rather than as independent responses to block-level cost shocks, all of which are consistent with these closures being strategic. The practice does not lower premiums for healthy shoppers: open deadpooling blocks are not priced below the same insurer's other blocks, and the cheapest quote in a market is no lower where deadpooling is more common.

Deadpooling's harm is large and concentrated. 17.6% of Medigap life-years through 2020 are spent in a plan that deadpools, and each year's closures let insurers charge an additional \$64.7 million over the following five years beyond what the higher claims of the remaining pool justify. Counting also the switching costs of enrollees who move and, using revealed preference, the consumer surplus lost by those driven out of coverage entirely, each dollar of this excess markup costs seniors \$1.67. Closure also lowers the share of premiums spent on care: the loss ratio of deadpooling blocks changes by -0.056 after closure, from a pre-closure mean of 0.795 (Table 2). Closed deadpooling blocks still clear Medigap's 0.65 floor, but they would fail the 0.85 standard that Medicare Advantage plans must meet.

Several states now give Medigap enrollees an annual right to switch insurers without underwriting, which in theory should end deadpooling. In practice, these rules help the enrollees who use them. Switching rules reduce the share of enrollment in closed deadpooling plans by 35% of its pre-adoption level, and once a rule is in force, enrollees in closed deadpooling blocks leave more after small and moderate premium increases. But the share of enrollment in deadpooling plans as a whole falls by an imprecisely estimated 18%, premiums and profits do not move significantly, and carriers do not exit in the years after adoption, although California and Oregon have fewer carriers in the long run. A switching right helps only the enrollees who shop, and most do not.

Our model shows why this matters for policy. A birthday rule reduces overall lost consumer surplus by only 5%, and raising consumer attention alone reduces it by 11%, but the two together reduce it by 83%: the switching right and consumer attention are complements. A direct ban on deadpooling barely changes lost surplus (a 0% reduction). It ends closures, but premiums in open blocks rise, because deadpooling parents no longer open successor blocks filled with younger, healthier enrollees. Among the rules a regulator can impose without relying on attention, automatic switching does the most, removing 53%.

These results speak to a broader tradeoff in insurance regulation. Limiting adverse selection has costs of its own (Handel, Hendel, and Whinston 2015), and lock-in is one way of limiting it. We document a cost of that solution: lock-in gives the insurer a beneficiary first enrolls with market power over them later, and corporate structure lets the parent company use that power while each of its plans complies with rules written for individual blocks. Where regulators rely on lock-in to stabilize risk pools, protecting consumers requires either reaching the parent company's pricing across its subsidiaries or making it easy for locked-in enrollees to leave.

Finally, deadpooling pushes some enrollees out of Medigap earlier than they would otherwise leave. If

they move to Medicare Advantage, the effect on government spending is unclear: Medigap raises Medicare spending through moral hazard (Cabral and Mahoney 2019), while Medicare Advantage raises it through higher payments (Medicare Payment Advisory Commission 2025). Which level of utilization is more efficient, and the health consequences of losing Medigap coverage at younger ages, remain questions for future work.

References

- Abaluck, Jason, and Jonathan Gruber. 2011. "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review* 101 (4): 1180–1210.
- . 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review* 106 (8): 2145–84.
- Ali, Riaz, Aimee Cicchiello, Morgan Hanger, Lesley Hellow, Ken Williams, and Gretchen Jacobson. 2021. "How Agents Influence Medicare Beneficiaries' Plan Choices."
- Anthem Blue Cross and Blue Shield. 2020. *Special Supplemental Benefits for the Chronically Ill*. Anthem Blue Cross and Blue Shield Medicare Advantage provider bulletin, January 2020. Accessed 23 October 2026.
- Armour, Philip, and Claire O'Hanlon. 2019. *How Does Supplemental Medicare Coverage Affect the Disabled Under-65 Population?: An Exploratory Analysis of the Health Effects of States' Medigap Policies for SSDI Beneficiaries*. Working Paper 25564. National Bureau of Economic Research, February.
- Atal, Juan Pablo. 2019. "Lock-in in Dynamic Health Insurance Contracts: Evidence from Chile." *SSRN Electronic Journal*.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance*." *The Quarterly Journal of Economics* 130 (4): 1623–1667.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Option." *The Quarterly Journal of Economics* 132 (3): 1319–1372.
- Borenstein, Severin, Jeffrey K. MacKie-Mason, and Janet S. Netz. 1995. "The Economics of Customer Lock-In and Market Power in Services." Chap. 8 in *The Service Productivity and Quality Challenge*, edited by Patrick T. Harker, 225–250. Kluwer Academic Publishers.
- Cabral, Marika. 2017. "Claim Timing and Ex Post Adverse Selection." *The Review of Economic Studies* 84 (1): 1–44.
- Cabral, Marika, and Neale Mahoney. 2019. "Externalities and Taxation of Supplemental Insurance: A Study of Medicare and Medigap." *American Economic Journal: Applied Economics* 11 (2): 37–73.
- Callaway, Brantly, and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics* 225 (2): 200–230.
- Centers for Medicare & Medicaid Services. 2023. *Medicare Current Beneficiary Survey, Survey File Public Use File, 2023*. Baltimore, MD: U.S. Department of Health and Human Services, Centers for Medicare & Medicaid Services.
- . 2026. *Joining a Plan*. Medicare.gov. Accessed 23 October 2026.
- Centers for Medicare and Medicaid Services. 2025. *NHE Fact Sheet*. National Health Expenditure Data.
- Choi, James J., David Laibson, and Brigitte C. Madrian. 2011. "\$100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans." *The Review of Economics and Statistics* 93 (3): 748–763.

- Claassen, Matthew. 2025. *Medicare Supplement Birthday Rule: a Disappointing Surprise?* Publication Title: MedigapSeminars, May.
- Conley, Timothy G., and Christopher R. Taber. 2011. "Inference with "Difference in Differences" with a Small Number of Policy Changes." *Review of Economics and Statistics* 93 (1): 113–125.
- Curto, Vilsa. 2023. "Pricing Regulations in Individual Health Insurance: Evidence from Medigap." *Journal of health economics* 91:102785.
- Cutler, David M., and Sarah J. Reber. 1998. "Paying for Health Insurance: The Trade-off between Competition and Adverse Selection." *The Quarterly Journal of Economics* 113 (2): 433–466.
- Decarolis, Francesco. 2015. "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" *American Economic Review* 105 (4): 1547–1580.
- Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives* 25 (1): 115–138.
- . 2023. "Empirical Analyses of Selection and Welfare in Insurance Markets: A Self-Indulgent Survey." *NBER Working paper* W31146.
- Ericson, Keith M. Marzilli. 2014. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy* 6 (1): 38–64.
- Fang, Hanming, Michael P. Keane, and Dan Silverman. 2008. "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market." *Journal of Political Economy* 116 (2): 303–350.
- Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." Chap. 31 in *Handbook of Industrial Organization*, edited by Mark Armstrong and Robert H. Porter, 3:1967–2072. Amsterdam: Elsevier/North-Holland.
- Finkelstein, Amy. 2004. "Minimum standards, insurance regulation and adverse selection: evidence from the Medigap market." *Journal of Public Economics* 88 (12): 2515–2547.
- Frank, Richard G., and Conrad Milhaupt. 2022. *Profits, medical loss ratios, and the ownership structure of Medicare Advantage plans*, July.
- Freed, Meredith, Alex Cottrill, Jeannie Fuglesten Biniek, and Tricia Neuman Published. 2023. *What Do People with Medicare Think About the Role of Marketing, Shopping for Medicare Options, and Their Coverage?* - Report - 10202, September.
- Freed, Meredith, Jeannie Fuglesten Biniek, Anthony Damico, and Tricia Neuman. 2026. *Medicare Advantage in 2026: Premiums, Out-of-Pocket Limits, Supplemental Benefits, and Prior Authorization*. KFF Issue Brief, June.
- Gabaix, X., and D. Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *The Quarterly Journal of Economics* 121 (2): 505–540.

- Ghili, Soheil, Benjamin Handel, Igal Hendel, and Michael D. Whinston. 2024. "Optimal Long-Term Health Insurance Contracts: Characterization, Computation, and Welfare Effects." *The Review of Economic Studies* 91 (2): 1085–1121.
- Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–277.
- Graham, Bryan S., Cristine Campos De Xavier Pinto, and Daniel Egel. 2012. "Inverse Probability Tilting for Moment Condition Models with Missing Data." *The Review of Economic Studies* 79 (3): 1053–1079.
- Handel, Benjamin, and Joshua Schwartzstein. 2018. "Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care?" *Journal of Economic Perspectives* 32 (1): 155–178.
- Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–2682.
- Handel, Benjamin R., Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica* 83 (4): 1261–1313.
- Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–2500.
- Heidhues, Paul, Mats Köster, and Botond Köszegi. 2023. "Steering Fallible Consumers." *The Economic Journal* 133 (652): 1430–1465.
- Hendel, Igal, and Alessandro Lizzeri. 2003. "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance." *The Quarterly Journal of Economics* 118 (1): 299–328.
- Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2017. "The impact of consumer inattention on insurer pricing in the Medicare Part D program." *The RAND Journal of Economics* 48 (4): 877–905.
- Insights, Medicare Market. 2025. *The Birthday Rule: Sept '25 Update on Medicare Supplement OE / GI Rules*. Publication Title: Medicare Market Insights, November.
- Ketcham, Jonathan D., Claudio Lucarelli, and Christopher A. Powers. 2015. "Paying Attention or Paying Too Much in Medicare Part D." *American Economic Review* 105 (1): 204–33.
- Klemperer, Paul. 1987. "Markets with Consumer Switching Costs." *The Quarterly Journal of Economics* 102 (2): 375–394.
- . 1995. "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade." *The Review of Economic Studies* 62 (4): 515–539.
- Koijen, Ralph S. J., and Motohiro Yogo. 2016. "Shadow Insurance." *Econometrica* 84 (3): 1265–1287.
- Lee, Sunmin, Ichiro Kawachi, Lisa F. Berkman, and Francine Grodstein. 2003. "Self-Rated Health and Mortality: Does the Relationship Extend to a Latino Population?" *American Journal of Epidemiology* 157 (12): 1042–1049.

- Lin, H., and M. R. Wildenbeest. 2020. "Nonparametric Estimation of Search Costs for Differentiated Products: Evidence from Medigap." *Journal of Business & Economic Statistics* 38 (4): 754–770.
- MacKinnon, James G., and Matthew D. Webb. 2017. "Wild Bootstrap Inference for Wildly Different Cluster Sizes." *Journal of Applied Econometrics* 32 (2): 233–254.
- Maestas, Nicole, Mathis Schroeder, and Dana Goldman. 2009. *Price Variation in Markets with Homogeneous Goods: The Case of Medigap*. Technical report w14679. Cambridge, MA: National Bureau of Economic Research, January.
- Medicare Payment Advisory Commission. 2025. *Report to the Congress: Medicare Payment Policy*. Washington, DC: Medicare Payment Advisory Commission, March.
- Mullainathan, Sendhil, Markus Nöth, and Antoinette Schoar. 2011. "The Market for Financial Advice: An Audit Study."
- NAIC. 2000. *NAIC Model Laws, Regulations, Guidelines and Other Resources: GUIDELINES FOR FILING OF RATES FOR INDIVIDUAL HEALTH INSURANCE FORMS*, July.
- National Association of Insurance Commissioners. 2022. *Annual Statement Instructions: Health*. Medicare Supplement Insurance Experience Exhibit, columns 14 and 18.
- . 2025. *2025 Survey of Medicare Supplement New or Innovative Benefits*. https://content.naic.org/sites/default/files/inline-files/2025NeworInnovativeBenefitChart_0.pdf. Model Regulation 651, Section 9.1 permits new or innovative benefits that do not alter standardized benefits and exclude outpatient prescription drugs.
- Ochieng, Nancy, Juliette Cubanski, Meredith Freed, and Tricia Neuman. 2024. *Nearly 7 in 10 Medicare Beneficiaries Did Not Compare Plans During Medicare's Open Enrollment Period*, September.
- Rothschild, Michael, and Joseph Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information*." *The Quarterly Journal of Economics* 90 (4): 629–649.
- Sant'Anna, Pedro H. C., and Jun Zhao. 2020. "Doubly Robust Difference-in-Differences Estimators." *Journal of Econometrics* 219 (1): 101–122.
- Social Security Administration. 2024. *Actuarial Life Table*. Office of the Chief Actuary, period life table for the Social Security area population, as used in the annual Trustees Report.
- Starc, Amanda. 2014. "Insurer pricing and consumer welfare: evidence from Medigap." *The RAND Journal of Economics* 45 (1): 198–220.

A Appendix: Data and Sample

A.1 The Standardized Medigap Plans

Figure A1 lists the benefits of each standardized Medigap plan letter.

Figure A1. Medigap Coverage Is Standardized into Ten Benefit Tiers

Medigap Benefit	Plan A	Plan B	Plan C	Plan D	Plan F*	Plan G*	Plan K	Plan L	Plan M	Plan N
Part A coinsurance and hospital costs up to an additional 365 days after Medicare benefits are used up	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Part B coinsurance or copayment	✓	✓	✓	✓	✓	✓	50%	75%	✓	✓***
Blood (first 3 pints)	✓	✓	✓	✓	✓	✓	50%	75%	✓	✓
Part A hospice care coinsurance or copayment	✓	✓	✓	✓	✓	✓	50%	75%	✓	✓
Skilled nursing facility care coinsurance	✗	✗	✓	✓	✓	✓	50%	75%	✓	✓
Part A deductible	✗	✓	✓	✓	✓	✓	50%	75%	50%	✓
Part B deductible	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
Part B excess charge	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗
Foreign travel exchange (up to plan limits)	✗	✗	80%	80%	80%	80%	✗	✗	80%	80%
Out-of-pocket limit**	N/A	N/A	N/A	N/A	N/A	N/A	\$7,060 in 2024	\$3,530 in 2024	N/A	N/A

Source: Centers for Medicare & Medicaid Services, *Choosing a Medigap Policy: A Guide to Health Insurance for People with Medicare* (CMS Product No. 02110), medicare.gov. Notes: The ten lettered plans are fixed by statute, so any two plans of the same letter cover identically regardless of insurer. Benefit design is therefore not a margin on which insurers compete, so price differences between plans of the same letter cannot be attributed to product differences.

A.2 Data Integrity

Plan Choice Data Plan choice data are created from the NAIC Medicare Supplement Insurance Experience Exhibit data, which are annual statements that insurance companies are required to file in most states under the NAIC Medicare Supplement Insurance Minimum Standards Model Regulation. Medicare Supplement policies are defined as all forms which are qualified under the Federal Certification Requirements or the NAIC Medicare Supplement Insurance Minimum Standards Model Act and Regulation. While there is a small possibility of filing mistakes or missing data, since the annual filings are required by law the NAIC data are the highest-quality data that exist and has been used in other research papers such as Maestas, Schroeder, and Goldman (2009) and Starc (2014). Cleaning of these data is deliberately light, because the filings are a statutory census rather than a sample and a dropped record is a lost block rather than a lost observation. We do not winsorize or trim the per-enrollee dollar outcomes used in the closure estimates; the one outcome we do winsorize is the loss ratio, at the first and ninety-ninth percentiles, because it is a ratio whose denominator can approach zero and whose raw upper tail therefore runs into the tens rather than reflecting anything about pricing. Blocks with no covered lives or no earned premium in a year cannot produce a per-enrollee outcome and are excluded from that year.

Panel Construction and Date Handling Dates are added to each year's data contemporaneously, so plans that close do not have closure dates in their earliest years in the panel. All dates are missing from the 2008

year of data. Date variables are assigned to the panel plan based on the maximum value observed in the data, since information like closure dates is often missing in early years when that information has not yet been determined. We link plans at the company $\times$ state $\times$ plan-letter $\times$ group-status level rather than the policy level, as many policies change their name over time. State restrictions on the number of plans a subsidiary can sell at once support this approach. Plans with nonconsecutive coverage years are dropped. This is a restriction on gaps in the *filings* and is distinct from the spell logic used to define deadpooling. The two are worth separating because they treat a gap in the data very differently. A *missing year* is an absence of filings: the block files in one year, does not file in the next, and files again later. That pattern is a data problem rather than a market event (we cannot tell whether the block existed in between) so those plans are dropped. A *spell* is the opposite case, where the block is present throughout but its behaviour indicates it was quietly reopened: enrollment in an already-closed block jumps by more than twenty percent and by at least fifty lives in a single year. Because an insurer can reopen a block without changing its NAIC code, treating such a series as one continuous block would average a closed period together with a fresh one. We therefore split it at each detected reopening and treat each spell as a distinct block, detecting reopenings iteratively until no further jumps are found. A genuine close-and-reopen cycle is thus retained rather than discarded, while a block with holes in its filing record is not. The restriction removes 4,196 plans, 17.7% of those in the collapsed panel, but only 4.2% of covered life-years, because the plans it drops are overwhelmingly small. Medicare SELECT plans are dropped. We also drop group plans, which are not available to all enrollees, except the plans one large carrier sells through a national seniors' membership organization.

CSG Data Cleaning Insurance companies report their rates directly to CSG, and the firm then independently verifies the pricing information. The data arrive as one file of rates per year. Each rate is posted for a rating area, a set of ZIP codes over which a carrier charges the same price, and CSG supplies a crosswalk from rating areas to ZIP codes, which we use to assign every rate to each ZIP code it covers. Within a plan letter, a quote varies across *individual pricing* cells, defined by the applicant characteristics on which a carrier is permitted to vary its price: age, gender, tobacco use, and whether the applicant qualifies for a household or couple discount. Each quote also records the plan's rating methodology (issue-age, attained-age, or community-rated). In this paper, we only consider non-tobacco rates, which understates what a smoker would be offered, and we use the rate without any household, couple or other discount. We also omit plans whose benefits differ from the standardized letter, such as versions that add dental, vision or fitness benefits or restrict the provider network. This leaves one quote per subsidiary $\times$ ZIP code $\times$ plan letter $\times$ age $\times$ gender $\times$ year. Where an analysis is at the state level, we average quotes across the ZIP codes a block is offered in within the state, and, where it combines sexes, we take the equal-weighted mean of the female and male quotes. Three further restrictions apply to the CSG data. We keep only the ten lettered plans fixed by statute, dropping pre-OBRA policies and the nonstandard plans sold in Massachusetts, Minnesota and Wisconsin, since those are not comparable across states and a plan letter does not mean the same thing in them. We drop records with no plan letter recorded. And we use the rate a plan posts for a standard applicant rather than any underwritten or rated-up figure, so a quoted price here is the price a clean applicant of that age would be charged, not the price someone with a flagged condition would be offered.

Parent-Subsidiary Mapping and Year Matching We use the NAIC *Listing of Companies* data to link companies to subsidiaries. Each block is assigned the parent recorded for its NAIC code in its own year to

account for acquisitions. Codes that match no parent are treated as their own parent.

A.3 Summary Statistics by Switching Regime

Table 1 pools all states. Table A1 reports the same rows separately for the three regulatory regimes of Figure 2: open enrollment, where a beneficiary outside the Medigap market can buy in at least annually; an annual switching window, where a beneficiary already holding a policy can move; and no protection, where leaving requires passing medical underwriting. The regime is a state-year attribute, so a state that adopts a rule mid-panel contributes its pre-adoption years to Panel C and its post-adoption years to Panel A or B. The panels partition the sample exactly: their plan-block-year counts sum to the 186,561 of the main table.

The point of the split is to show that the deadpooling pattern is not an artifact of the locked-in states, and that the pooled table is nonetheless dominated by them. The ordering that Section 4.7 draws from the pooled table holds in every panel: deadpooling blocks charge less than blocks that never close while they are open, charge more than both non-deadpooling blocks and blocks that never close once they close, and lose most of their enrollment in between. What differs across panels is scale, not sign. The no-protection states supply 164,316 of the 186,561 block-years and the open-enrollment states only 4,674, so the pooled means in Table 1 sit close to Panel C throughout. The open-enrollment states are also a different kind of market (354 blocks holding 1,977 enrollees on average, against 15,687 blocks averaging 757 where there is no protection) which is why their much higher premium levels move the pooled figures so little.

B Appendix: Measuring Underwriting Eligibility from the MCBS

Two claims in the main text rest on an exercise that matches the health conditions recorded in the Medicare Current Beneficiary Survey (MCBS) to the medical-underwriting criteria Medigap insurers actually apply: (i) how many Medicare beneficiaries carry a condition that would draw underwriting scrutiny, and (ii) that 26.8% of current Medigap enrollees could pass underwriting and switch to another Medigap while 73.2% could not (Sections 2.2 and 7.3). This appendix documents how each number is built.

Who carries an underwriting condition. Prevalence differs by how a beneficiary is covered, and the paper’s headline figure is the market-wide one. Across all Medicare beneficiaries, 69.5% carry at least one condition an underwriter would act on under the conservative mapping (88.9% under the broad one). Among Medigap enrollees the share is 66.8% (87.8% broad), and among Medicare Advantage enrollees 70.5% (90.0% broad). That Medigap enrollees are the *least* likely of the three groups to carry such a condition is what the underwriting screen is for: it is the population that has already been selected on health at entry.

MCBS data. We stack the MCBS fall waves for 2017–2023 and survey-weight throughout using the public-use weight PUFFWGT. The MCBS public-use file does not flag Medigap directly, so we proxy it by traditional-Medicare enrollees who also report self-purchased private insurance (INS_D_MADV=2 and INS_D_PVSELF=1); this is the best available proxy. The underwriting criteria come from the CSG Actuarial 2022 “Supp Only” data feed, which records, for each insurer’s Medigap product in each state, the health conditions it screens and the lookback window it uses; a blank criterion means the condition is not used in underwriting for that product.

Table A1. Summary Statistics by State Switching Regime

	Means by block status				
	Full sample	Open, DP	Closed, DP	Non-DP, all	Never close
Panel A: Open enrollment (annual guaranteed entry to the market)					
<i>CSG: Quoted rates, 2014–2025</i>					
Annual rate, age 65	\$3,360	—	\$3,898	\$3,359	\$3,336
Annual rate, age 70	\$3,360	—	\$3,898	\$3,359	\$3,336
Annual rate, age 75	\$3,360	—	\$3,898	\$3,359	\$3,336
Annual rate, age 80	\$3,360	—	\$3,898	\$3,359	\$3,336
Annual rate, age 85	\$3,360	—	\$3,898	\$3,359	\$3,336
<i>NAIC: Collected premiums, 2006–2025</i>					
Mean per-person profit	\$685	\$894	\$1,094	\$682	\$678
Mean loss ratio	0.80	0.74	0.77	0.80	0.80
Mean per-person premium	\$3,376	\$3,277	\$3,964	\$3,373	\$3,350
Mean per-person claims	\$2,691	\$2,383	\$2,870	\$2,691	\$2,672
Mean enrollment	1,977	458	60	2,479	4,242
Plan-block-years	4,674	73	910	3,691	1,974
Unique blocks	354	28	82	272	136
Parent companies	54	11	20	50	27
Panel B: Annual switching window					
<i>CSG: Quoted rates, 2014–2025</i>					
Annual rate, age 65	\$2,582	\$2,722	\$3,097	\$2,563	\$2,422
Annual rate, age 70	\$2,913	\$3,054	\$3,483	\$2,892	\$2,728
Annual rate, age 75	\$3,365	\$3,559	\$4,076	\$3,338	\$3,097
Annual rate, age 80	\$3,652	\$4,066	\$4,614	\$3,612	\$3,396
Annual rate, age 85	\$3,862	\$4,622	\$5,153	\$3,807	\$3,601
<i>NAIC: Collected premiums, 2006–2025</i>					
Mean per-person profit	\$617	\$573	\$1,036	\$597	\$512
Mean loss ratio	0.80	0.78	0.75	0.80	0.82
Mean per-person premium	\$2,965	\$2,526	\$3,854	\$2,934	\$2,741
Mean per-person claims	\$2,348	\$1,954	\$2,819	\$2,337	\$2,229
Mean enrollment	1,029	972	251	1,223	1,483
Plan-block-years	17,571	535	3,372	13,664	7,355
Unique blocks	2,324	168	525	1,799	1,053
Parent companies	141	40	62	135	80
Panel C: No switching protection					
<i>CSG: Quoted rates, 2014–2025</i>					
Annual rate, age 65	\$2,544	\$2,539	\$2,792	\$2,516	\$2,493
Annual rate, age 70	\$3,017	\$2,940	\$3,206	\$3,000	\$2,967
Annual rate, age 75	\$3,415	\$3,353	\$3,710	\$3,386	\$3,313
Annual rate, age 80	\$3,777	\$3,706	\$4,149	\$3,739	\$3,622
Annual rate, age 85	\$4,028	\$4,026	\$4,522	\$3,972	\$3,859
<i>NAIC: Collected premiums, 2006–2025</i>					
Mean per-person profit	\$622	\$631	\$869	\$593	\$533
Mean loss ratio	0.80	0.77	0.78	0.81	0.82
Mean per-person premium	\$2,948	\$2,499	\$3,480	\$2,921	\$2,813
Mean per-person claims	\$2,326	\$1,868	\$2,611	\$2,329	\$2,281
Mean enrollment	757	713	354	873	1,293
Plan-block-years	164,316	11,159	33,432	119,725	59,374
Unique blocks	15,687	2,994	4,414	11,245	6,032
Parent companies	220	104	119	214	134

Notes: Same panel, weighting, deflator, CSG construction and five-year lookforward cutoff as Table 1. The columns are the full sample; deadpooling blocks before closure (“Open, DP”) and after it (“Closed, DP”); all non-deadpooling blocks (“Non-DP, all”), a column Table 1 does not report; and blocks that never close (“Never close”), defined as in Table 1; the three panels differ only in which state-years they retain, and between them they exhaust the sample. Panel C is the sample on which every mechanism test in the paper is run. Premiums, claims and profits are per enrollee and enrollment-weighted, dollar figures are in January 2025 dollars, and the loss ratio is total claims divided by total premiums. CSG rows are the equal-weight mean of the male and female non-smoker quote, 2014–2025 only, with closed blocks that are no longer quoted priced at their renewal schedule. A dash marks a cell with no CSG price. The open-enrollment states community-rate Medigap, so their quote does not vary with age.

CSG data: what a criteria record contains. The underwriting feed is organised as one record per selling entity $\times$ state $\times$ product, and 161 NAIC entities supply rules. Each record lists the health conditions that entity screens for in that market and, for each, the lookback window over which the condition is asked: whether it is a current diagnosis, a treatment or an episode within some number of years, or a scheduled procedure. A blank cell means the condition is not used in underwriting for that product, which is why the file has to be read as a grid rather than as a list of conditions—absence is informative. Two features of the data matter for how it is used here. First, the conditions are not free text: after harmonising wording across carriers they reduce to 16 distinct condition families, which is what makes a crosswalk to survey-reported diagnoses feasible at all. Second, carriers differ far less than the number of records suggests: across 2,426 entity $\times$ state markets there are only 56 distinct rule sets, so most of the market underwrites on one of a small number of common templates.

CSG data: coverage of the market. The CSG feed is close to a census of the market rather than a sample of it. The 2022 underwriting guides we use come from companies that, in the NAIC panel that same year, account for 91.8% of all Medigap covered lives at the selling-entity (NAIC) level and 96.8% at the parent-company level; total 2022 Medigap enrollment is 11,325,615. In particular, we have underwriting rules for 10 of the 10 largest parent companies, which together hold 75.7% of the market. The companies for which we lack rules are small carriers; nothing about the exercise turns on them.


Mapping conditions. We translate CSG’s underwriting condition strings into 16 indicators the MCBS can measure, listed in Table A2. Because the MCBS records “ever” diagnoses while most underwriting criteria use one- or two-year lookbacks, these indicators are an *upper* bound on the conditions a given insurer would actually act on. Conversely, conditions the MCBS cannot see (multiple sclerosis, Parkinson’s, ALS, organ transplant, cirrhosis, HIV/AIDS, hepatitis, alcoholism, lupus, and most surgical history) are dropped, so the resulting rejection rates are a *lower* bound along that dimension. In the circulatory rows we exclude condition strings beginning “Diabetes With...” so that diabetes-complication entries are counted once, through the diabetes flags, not double-counted as heart or artery disease.

Prevalence of underwriting-relevant conditions (the ~70% figure). We form two composites. The *conservative* composite flags a beneficiary who has at least one of nine CSG-mapped primary categories (circulatory disease, internal cancer, insulin-dependent diabetes, respiratory disease, severe musculoskeletal disease, nervous-system disease, psychosis, kidney disease, or Medicare eligibility through disability). The *broad* composite adds hypertension, any diabetes, and depression, which some insurers flag even without complications. Pooling 2017–2023, 69.5% of all Medicare beneficiaries carry a conservative-composite condition (88.9% under the broad composite); among Medigap holders the figures are 66.8% and 87.8%. Because the “ever”-diagnosis measure overstates what a lookback-limited underwriter would act on, we describe this in the text as “about 70%” who would not be offered a plan, or offered one only at elevated rates.

Market-specific rejection and the switch-versus-exit split (the 26.8%/73.2% figures). For each of the 2,426 company $\times$ state underwriting guides in the 2022 CSG feed we record which of the 16 conditions it declines on; collapsing these to their distinct patterns yields 56 unique rejection “fingerprints.” For every MCBS respondent we then compute *rej_rate*, the share of all 2,426 guides that would decline them given their

conditions (each company×state guide counting equally). Figure A2 plots that distribution in full, and its *shape*, not any single summary statistic, is what carries the argument. It is starkly bimodal with a hollow interior: enrollees pile up either at zero (accepted everywhere) or against the top of the range, and the intermediate bins are essentially unpopulated. 26.8% of Medigap holders would be accepted in *every* market ($\text{rej_rate} = 0$); at the opposite pole most of the rest would be declined in nearly every market, and the weighted mean, 68.5%, falls in a region where almost nobody actually sits. We therefore read the 26.8% accepted-everywhere group as the enrollees who can obtain another Medigap when their plan deadpools, and the complementary 73.2% as those who cannot. The bimodality is what makes this usable: the split is nearly identical whether “can switch” requires acceptance in every market or in merely one.

Figure A2. Distribution of the Share of Medigap Markets That Would Decline Each Enrollee



Notes. Each MCBS respondent holding Medigap is scored against every one of the 2,426 company×state underwriting guides in the 2022 CSG feed; rej_rate is the share of those guides that would decline them given their recorded conditions. Bars are survey-weighted and five percentage points wide. The sample is the MCBS fall waves pooled over 2017–2023, crossed with the 2022 CSG markets. The distribution is bimodal with an empty middle: the modal enrollee is either accepted everywhere or rejected almost everywhere, and virtually no one occupies the intermediate range. This is what licenses reading the accepted-everywhere group as the enrollees who can obtain replacement coverage when their block deadpools, and the remainder as those who cannot—the split is insensitive to where in the empty interior one draws the line.

Limitations. Five limitations bound these figures, and their directions are known. (1) The “ever”-diagnosis measure overstates the conditions an underwriter with a finite lookback would act on, so the prevalence figures are upper bounds. We can now bound the error from the other side as well. The CSG criteria field records the lookback window each screen applies, and only 26.7% of screening rows use an “ever” horizon; the majority specify a finite window, most commonly two years. Recomputing the market-rejection calculation while crediting a market with screening a condition *only* where it does so on an “ever” lookback (the horizon to which the MCBS measure actually corresponds, and holding all 2,426 markets in the denominator) gives a mean rejection rate of 29.4%, against the headline 68.5%. That figure is a lower bound in precisely the way the headline is an upper bound: it credits no rejection whatever to a market screening on a two-year window, although such a market does decline recently diagnosed applicants. The truth lies between them. What

matters for our purposes is that the quantity the switch-versus-exit split actually rests on (the share accepted in *every* market) is unchanged between the two calculations (26.8% against 26.8%), because an enrollee carrying no screenable condition passes irrespective of the lookback horizon. The harm estimates in Section 7.3 are therefore insensitive to this bound. (2) Conditions the MCBS cannot observe are omitted, so the market-rejection rates are lower bounds. (3) The Medigap proxy may absorb a few non-Medigap self-purchased plans. (4) HLT_OCKIDNY (kidney disease) is unavailable in 2017 and HLT_OCOSARTH (osteoarthritis) in 2019, and are treated as missing those years. (5) Standard errors would be understated were we to report them, because the MCBS public-use file ships a single analytic weight (puf_wgt) and no replicate weights; the replicate weights needed for correct variance estimation are available only in the restricted-use file, which we do not hold. We therefore report prevalences as point estimates and do not attach standard errors to them. Finally, passing underwriting is necessary but not sufficient for a beneficiary to actually switch (the NAIC data cannot observe destinations) so the 26.8% switcher share is if anything too high, which *understates* the destroyed surplus attributed to full exiters in Section 7.3.

Table A2. Sixteen Condition Groups Link CSG Screening Rules to MCBS Conditions

Condition group	CSG 2022 condition strings matched	MCBS variable
Stroke	"Stroke"; "TIA"	HLT_OCSTROKE
Heart attack	"Heart Attack"	HLT_OCMYOCAR
Congestive heart failure	"Congestive Heart Failure"	HLT_OCCFAIL
Coronary / artery disease	"Coronary"; "Artery Disease"; "Artery Blockage"; "Angina"; "Carotid"	HLT_OCCHD
Other heart condition	"Heart Rhythm"; "Cardiomyopathy"; "Enlarged Heart"; "Fibrillation"; "Angioplasty"; "Any Heart Condition"; "Valve Surgery"; "Bypass"; "Pacemaker"	HLT_OCHRTCND
Peripheral / arterial disease	"Peripheral Vascular"; "Arterial Disease"	HLT_OCARTERY
Internal cancer	"Cancer"; "Leukemia"; "Melanoma"; "Lymphoma"; "Myeloma"; "Hodgkin"; "Bone Marrow"; "Stem Cell"	HLT_OCCANCER
Insulin-dependent diabetes	"Diabetes Dependent on Insulin"	HLT_DIAINSUL
Other diabetes	category "Diabetes" (excluding insulin-dependent)	HLT_OCBETES
COPD / respiratory	"COPD"; "Emphysema"; "Pulmonary"; "Lung"	HLT_OCEMPHYS
Rheumatoid arthritis	"Rheumatoid Arthritis"	HLT_OCARTHRI
Osteoporosis with fracture	"Osteoporosis With Fractures"	HLT_OCOSTEOP & HLT_OCBRKHIP
Alzheimer's / dementia	"Alzheimer"; "Dementia"; "Cognitive Disorder"	HLT_ALZDEM
Psychiatric / organic brain	"Organic Brain"; "Mental or Nervous"; "Psychiatric"	HLT_OCPSYCHO
Kidney disease	"Kidney"; "Renal"	HLT_OCKIDNY (2018+)
Disability / functional	"Bedridden"; "Wheelchair"; "Nursing Facility"; "Activities of Daily Living"; "Home Health"	disabled (ADM_H_MEDSTA=2)
No MCBS counterpart (dropped)	Multiple sclerosis; Parkinson's; ALS; organ transplant; cirrhosis; HIV/AIDS; hepatitis; alcoholism; lupus; most surgical history	—

Notes: Each row is one underwriting flag used in both the prevalence composites and the market-rejection exercise. CSG condition strings are matched by substring within the 2022 "Supp Only" feed; a market screens a flag if any listed string appears with a non-empty lookback criterion. "Osteoporosis with fracture" requires both the osteoporosis and the broken-hip flags. Circulatory rows exclude strings beginning "Diabetes With..." so diabetes complications are counted only through the diabetes flags. Conditions in the CSG feed with no MCBS counterpart (e.g. multiple sclerosis, Parkinson's, cirrhosis, organ transplant) are omitted. The crosswalk is deliberately asymmetric in its errors: because the MCBS records "ever" diagnoses against underwriters' one- or two-year lookbacks, the prevalence figures are upper bounds, while dropping conditions the MCBS cannot see (multiple sclerosis, Parkinson's, transplant, cirrhosis, HIV) makes the rejection rates lower bounds.

Cross-check against the agent survey. The agent survey of Appendix F asks brokers to divide their own clients on a closed block by what those clients were actually offered, and 23.6% were offered underwritten Medigap at a competitive rate—against the 23.6% implied by MCBS condition prevalence. One figure comes from clinical prevalence applied to observed retention, the other from brokers describing books they placed themselves. We do not read the agreement as more precise than the inputs allow, but it is reassuring for a split the rest of this accounting rests on.

C Appendix: Robustness of the Closure Estimates

This appendix collects robustness checks on Table 2. The main estimates pool all standardized plan letters and weight by enrollment.

The modal closure cohort. Table A3 re-estimates the deadpooling column of Table 2 with and without the 2010 closure cohort. Column 1 reports the specification from Table 2's. Dropping 2010, the modal closure year, lowers the premium ATT from \$1,005 to \$842, claims from \$606 to \$512, profit from \$398 to \$330, and the loss-ratio effect from -0.056 to -0.047 , and every estimate remains significant at the one percent level. The effect is somewhat smaller without 2010, but it still large and precisely estimated. 2010 closures are not observed in CSG, so we do not show reestimated premium results.

Table A3. The Closure Effect Does Not Rest on the 2010 Closure Cohort

	(1) Baseline	(2) Drop 2010 closers	Pre-Closure Means
<i>Panel A: Per-Person Premiums (2025 dollars)</i>			
Closure, deadpooling plans	\$1,005*** (235)	\$842*** (222)	\$2,347
<i>Panel B: Per-Person Claims (2025 dollars)</i>			
Closure, deadpooling plans	\$606*** (139)	\$512*** (130)	\$1,821
<i>Panel C: Per-Person Profits (2025 dollars)</i>			
Closure, deadpooling plans	\$398*** (100)	\$330*** (96)	\$526
<i>Panel D: Loss Ratio (claims / premiums)</i>			
Closure, deadpooling plans	-0.056*** (0.018)	-0.047*** (0.016)	0.795
State controls	X	X	
Plan type controls	X	X	
N (plan-block-years, deadpooling only)	111,097	99,573	
Unique blocks	12,478	11,303	
Subsidiaries	291	286	
Parent companies	207	203	

Notes: Callaway–Sant’Anna (2021) overall ATT of closure on each per-person outcome, deadpooling closures against all never-closing blocks, with state and plan-type controls by inverse probability tilting (`drip`), a universal base period (`long2`), enrollment weights and standard errors clustered at the NAIC insurer level, exactly as in column 3 of Table 2. Column 1 is that column; column 2 drops every block that closed in 2010, leaving the never-closing controls unchanged. The right-hand column is the lives-weighted pre-closure mean of the outcome among deadpooling closers. N is the number of plan-block-years used in estimation; unique blocks, subsidiaries and parent companies are counted on each column’s sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Imputed closure dates. A block’s closure year is read off the NAIC filings when possible and imputed otherwise. Blocks that stops filing before 2025 without a recorded closure are treated as closing in their last filing year. Among deadpooling closures, imputed closures are 15.6% of closures but 3.5% of enrollment in the year before closure. We re-estimate the event study of Figure 3 on recorded closers alone, dropping every

block whose closure year is imputed, which removes 2.7% of the panel ($N = 111,097$ to $N = 108,136$). Our results are robust to this specification.

C.1 Pre-Trend Tests

Identification in Table 2 rests on parallel trends, and this section reports the test of it. For each outcome we run a joint Wald test that all five pre-treatment event-time effects ($t = -6$ to -2 ; $t = -1$ is the base period) are zero, following equation (4), so a *large* p -value is the favourable outcome. The test is run on event-time aggregates rather than on the individual cohort-period cells, whose cluster-robust covariance is rank-deficient here (Section 5). We test both arms: the deadpooling plans of the main table, and the non-deadpooling closers of Appendix C.3. Column numbers below refer to the three specifications of Table 2, and column 3 is the fully adjusted one.

The two arms behave differently under the test, and the difference governs how we read the estimates. For non-deadpooling closers, every outcome passes comfortably in Column 3 (0.104 for premiums, 0.341 for claims, 0.180 for profit, 0.452 for the loss ratio). For deadpooling closers the picture is mixed, and it splits by outcome rather than uniformly. Premiums and claims reject at the five percent level in *every* column, Column 3 included (0.004 and 0.025), and the covariate adjustments do not rescue them. Profit and the loss ratio, by contrast, fail to reject in every column (0.188 and 0.260 in Column 3). Under `dripw` the premium test also rejects in columns 1 and 2 ($p = 0.010$ and 0.039) but not in column 3 ($p = 0.067$). Appendix C.5 explains why that column-3 pass reflects a loss of power and is not evidence of cleaner identification.

C.2 Predictors of Treatment Timing

Table A4 reports the tests discussed in Sections 5.2 and 5.4: whether the characteristics of a block predict when it deadpools (Panel A), and whether the characteristics of a state’s Medigap market predict when the state adopts a switching law (Panel B). Each cell comes from a separate regression on one characteristic.

C.3 Closure Estimates for Deadpooling and Non-Deadpooling Plans Together

Table 2 in the main text reports the closure estimates for deadpooling plans only. This appendix reproduces the same specification with both arms shown together: the deadpooling plans and, as a separate row in each panel, the disjoint set of plans that close without their parent reopening. The two groups are mutually exclusive, so the rows can be read against one another directly.

The contrast is sharpest in columns 1 and 2, where both arms are precisely estimated and the deadpooling effect is two to three times larger. In column 3 the non-deadpooling arm loses most of its precision, and the reason is mechanical rather than substantive: adding the plan-type adjustment to the propensity model drops every cohort-year cell before 2011 on that arm, which removes the 2010 closure wave—the largest in the panel. The estimates in that column should be read with that in mind.

C.4 Clustering at the Parent Company

The main closure estimates cluster standard errors at the NAIC insurer (Section 5). Many insurers belong to a larger parent group, and a parent can set pricing and closure policy across all of its subsidiaries at once, so

Table A4. Market Conditions Do Not Predict Switching Laws or Deadpooling Closures

<i>Panel A. Deadpooling closure</i>				
	Unweighted		Lives-weighted	
	(1) Market FE	(2) + Sub. × year FE	(3) Market FE	(4) + Sub. × year FE
Log claims per enrollee	−0.21 (0.19)	−0.21** (0.10)	−2.97* (1.80)	2.66** (1.09)
Claims growth	−0.10 (0.07)	−0.04 (0.04)	2.41** (1.22)	−1.17 (0.75)
Log premium per enrollee	−0.30 (0.31)	−0.21 (0.20)	−1.65 (1.05)	1.41** (0.60)
Premium growth	−0.05 (0.14)	−0.05 (0.08)	1.53** (0.62)	−0.45 (0.40)
Loss ratio	0.03 (0.06)	0.05 (0.05)	0.40 (2.16)	2.34** (1.16)
Log enrollment	0.04 (0.44)	−0.00 (0.21)	−1.86** (0.75)	2.31*** (0.89)
Enrollment growth	−0.21 (0.17)	−0.24*** (0.09)	−0.05 (0.36)	−0.98*** (0.32)
Block age	−0.77** (0.32)	−0.58** (0.24)	−1.79*** (0.64)	−0.23 (0.39)
Log quoted price at 65 (CSG)	1.01 (0.90)	−0.05 (0.30)	−0.96 (0.72)	−0.53 (0.84)
Quoted-price growth (CSG)	−0.09 (0.30)	−0.02 (0.15)	−0.36 (0.35)	0.12 (0.37)
Attained-age rating (0/1)	2.25*** (0.77)	0.32 (0.58)	4.65*** (1.08)	−0.13 (1.52)
<i>Joint tests (p-value of F-test that all coefficients are zero)</i>				
Claims, premiums, loss ratio, enrollment, age	0.502	0.167	0.012	0.156
Prices (premium and quote, level and growth)	0.713	0.304	0.152	0.147
All characteristics	0.119	0.369	0.000	0.134
Baseline hazard (%)			3.2	
Deadpooling closures			3,613	
Block-years			112,242	
<i>Panel B. Switching-law adoption</i>				
	Hazard, pp per SD		2022 adopters – never	
	(1) Unweighted	(2) Lives-weighted	(3) Unweighted	(4) Lives-weighted
Premium per enrollee	0.20 (0.29)	0.32 (0.38)	174*** (59)	171*** (51)
Claims per enrollee	0.13 (0.23)	0.28 (0.46)	89* (45)	111** (49)
Loss ratio	−0.07 (0.13)	−0.08 (0.09)	−0.014* (0.008)	−0.007 (0.007)
Premium growth	−0.06 (0.14)	−0.15 (0.20)	−0.003 (0.004)	−0.001 (0.002)
Claims growth	−0.14 (0.21)	0.05 (0.23)	−0.002 (0.006)	−0.005 (0.003)
Log enrollment	0.20 (0.24)	−0.01 (0.70)		
Enrollment growth	0.30 (0.23)	0.48 (0.36)	−0.001 (0.006)	−0.002 (0.004)
Number of parent companies	0.12 (0.25)	0.21 (0.82)	3.3 (6.1)	3.8 (5.8)
Parent entries	−0.20 (0.28)	−0.22 (0.39)	−0.31 (0.35)	−0.54** (0.25)
Parent exits	−0.13 (0.28)	−0.72 (0.52)	−0.16 (0.16)	−0.36** (0.15)
New-block entry rate	−0.07 (0.15)	0.29 (0.47)	−0.006 (0.005)	−0.002 (0.003)
Closing share of lives	0.04 (0.22)	0.54 (0.65)		
Deadpooling share of lives	0.51 (0.39)	0.20 (0.37)	0.054 (0.038)	0.001 (0.028)
<i>Joint tests (p-value of F-test that all coefficients are zero)</i>				
All characteristics	0.938	0.903	0.898	0.946
Premium and claims, level and growth	0.887	0.788	0.398	0.751
Adoptions		6		4
Baseline hazard (%)		0.8		
State-years		734		123

Notes: Coefficients are *not* jointly estimated: each coefficient and standard error comes from a separate regression containing only that one characteristic and the fixed effects. The “Joint tests” rows instead include all the listed characteristics together in one regression and report the *p*-value of an *F*-test that all of their coefficients are zero. Panel A estimates $D_{i,t+1} = \beta x_{it} + \alpha_{s,t} + \varepsilon_{it}$ over open block-years, where $D_{i,t+1}$ indicates that block *i* closes in *t* + 1 and is a deadpooling closure, x_{it} is the block’s characteristic in year *t*, standardized, and $\alpha_{s,t}$ are state × plan-letter × year fixed effects; columns (2) and (4) add subsidiary × year fixed effects γ_{mt} . Coefficients are percentage points per standard deviation, except attained-age rating (0/1), and standard errors are clustered by NAIC company. Columns (3) and (4) weight by the block’s covered lives in *t*. Quoted prices are CSG non-tobacco quotes at age 65, averaged across sexes and ZIP codes within a state, available from 2014, so rows and joint tests using them have smaller samples. Panel B, columns (1) and (2), estimate $A_{st} = \beta x_{s,t-1} + \delta_t + u_{st}$ over state-years that have not yet adopted, where A_{st} indicates that state *s* adopts a switching law in year *t* and δ_t are year fixed effects; coefficients are percentage points per standard deviation, standard errors are clustered by state, and column (2) weights by the state’s covered lives in *t* − 1. Columns (3) and (4) estimate $x_{st} = \beta \mathbf{1}\{s \text{ adopted in 2022}\} + u_{st}$ over 2019–2021 for the four 2022 adopters and never-adopting states, in natural units (dollars for premiums and claims, log points for growth rates, shares for rates); column (4) weights by covered lives. States with switching or open-enrollment rules that predate the panel are excluded, as in the switching-rule analysis. Growth rates are year-on-year log changes. **p* < .1, ***p* < .05, ****p* < .01.

Table A5. Closure Estimates, Deadpooling and Non-Deadpooling Plans

	(1)	(2)	(3)
<i>Panel A: Per-Person Premiums (2025 \$)</i>			
Closure, Deadpooling Plans Only	\$1,116*** (242)	\$1,049*** (236)	\$1,005*** (235)
Closure, Non-Deadpooling	\$341** (139)	\$344** (150)	\$329** (161)
Pre-closure mean: \$2,347 (deadpooling); \$2,746 (non-DP)			
Pre-trend holds, deadpooling			
Pre-trend holds, non-deadpooling		X	X
<i>Panel B: Per-Person Claims (2025 \$)</i>			
Closure, Deadpooling Plans Only	\$724*** (139)	\$676*** (146)	\$606*** (139)
Closure, Non-Deadpooling	\$195** (79)	\$167** (81)	\$146 (96)
Pre-closure mean: \$1,821 (deadpooling); \$2,161 (non-DP)			
Pre-trend holds, deadpooling			
Pre-trend holds, non-deadpooling		X	X
<i>Panel C: Per-Person Profits (2025 \$)</i>			
Closure, Deadpooling Plans Only	\$392*** (107)	\$373*** (96)	\$398*** (100)
Closure, Non-Deadpooling	\$147** (70)	\$177** (75)	\$183*** (71)
Pre-closure mean: \$526 (deadpooling); \$585 (non-DP)			
Pre-trend holds, deadpooling	X	X	X
Pre-trend holds, non-deadpooling		X	X
<i>Panel D: Loss Ratio (claims / premiums)</i>			
Closure, Deadpooling Plans Only	-0.035* (0.019)	-0.038** (0.015)	-0.056*** (0.018)
Closure, Non-Deadpooling	-0.006 (0.019)	-0.026** (0.012)	-0.029*** (0.010)
Pre-closure mean: 0.795 (deadpooling); 0.796 (non-DP)			
Pre-trend holds, deadpooling	X	X	X
Pre-trend holds, non-deadpooling	X	X	X
State controls		X	X
Plan type controls			X
N (plan-block-years, deadpooling only)	111,097	111,097	111,097
N (plan-block-years, non-deadpooling)	150,222	150,222	150,222
Unique blocks (deadpooling / non-deadpooling)	12,478 / 14,692		
Subsidiaries (deadpooling / non-deadpooling)	291 / 322		
Parent companies (deadpooling / non-deadpooling)	207 / 236		

Notes: Specification, sample, weighting, clustering, estimator (drimp) and pre-trend tests are identical to Table 2; see that table's note. The only difference is that both arms are displayed. "Closure, Deadpooling Plans Only" is the row reproduced in the main text. "Closure, Non-Deadpooling" restricts to plans that close and whose parent does not open a successor in the same market within five years; deadpooling closers are dropped from that arm's sample entirely rather than left in as controls, so the two rows are disjoint groups rather than a subset and its superset. Standard errors clustered at the NAIC insurer level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

shocks may be correlated across sibling insurers as well as within one. This appendix re-estimates the closure exhibits with standard errors clustered at the parent company instead.

Defining the parent. Parent assignments change over the panel as insurers are bought and sold, so many blocks carry more than one parent code over their lives. The Callaway–Sant’Anna estimator requires every block to sit in a single cluster, so we assign each block its *modal* parent: the parent group under which it carried the most covered lives over its whole life. Everything else—sample, weighting, covariates, estimator (drimp) and universal base period—is identical to Table 2, so every point estimate is unchanged and only the standard errors, significance stars and pre-trend tests can move.

Results. For deadpooling closers (Table A6) every estimate remains significant at the one percent level for premiums, claims and profit in every column, and the loss-ratio estimates become more precise rather than less.

Why coarser clusters give smaller standard errors. Clustering on a coarser unit can shrink a standard error only if the contributions of sibling insurers to the estimate are negatively correlated. Decomposing the estimator’s influence function by insurer shows that the shrinkage in the deadpooling arm comes entirely from a single large parent group, whose subsidiaries’ contributions largely offset one another. Most of the offset is between the deadpooling closures of that group’s different subsidiaries; a smaller part is between one subsidiary’s closures and a sibling subsidiary’s never-closing blocks, which is the pattern that closing a block in one subsidiary while selling a successor in another would produce. Across every other parent the cross-subsidiary terms are small and slightly positive, so without that group parent clustering would widen the intervals a little, as one would expect. The parent-clustered standard errors for this arm therefore rest heavily on one group’s internal covariance, which is a reason to treat them as a check rather than as a more conservative replacement for the NAIC-clustered errors in the main text.

Few effective clusters. The same concentration limits what either clustering can deliver. Because covered lives are concentrated in a handful of parents, a single cluster accounts for most of the variance of the estimate under both NAIC and parent clustering, so the effective number of clusters is far below the nominal count, and in the premium specification it is close to two. Asymptotic cluster-robust inference, and small-sample corrections such as the wild cluster bootstrap, are unreliable at that scale.

Pre-trends. One verdict changes. Under parent clustering the joint pre-trend test for deadpooling closers’ profit rejects at the five percent level in Column 3, and the loss ratio’s Column 3 test passes only narrowly. Deadpooling closers’ profit therefore passes the pre-trend test in every column of Table 2, but not in the fully adjusted column once shocks are allowed to be shared across a parent’s subsidiaries; the loss ratio passes in every column under either clustering.

C.5 Estimator Choice

The closure estimates in Section 6.1 use `csdid`’s `drimp` estimator rather than its default, `dripw`. This appendix explains the difference and compares the premium and profit effects under the two. The headline is that our findings are robust to the choice (the sign and order of magnitude of every estimate survive) but we prefer `drimp` because `dripw` discards a large share of the sample.

Table A6. Table 2 with Parent-Company-Clustered Standard Errors

	(1)	(2)	(3)	Pre-Closure Means
<i>Panel A: Per-Person Premiums (2025 \$)</i>				
Closure (deadpooling plans)	\$1,116*** (187)	\$1,049*** (191)	\$1,005*** (198)	\$2,347
<i>Panel B: Per-Person Claims (2025 \$)</i>				
Closure (deadpooling plans)	\$724*** (106)	\$676*** (117)	\$606*** (125)	\$1,821
<i>Panel C: Per-Person Profits (2025 \$)</i>				
Closure (deadpooling plans)	\$392*** (87)	\$373*** (83)	\$398*** (77)	\$526
<i>Panel D: Loss Ratio (claims / premiums)</i>				
Closure (deadpooling plans)	-0.035** (0.016)	-0.038*** (0.013)	-0.056*** (0.011)	0.795
State controls		X	X	
Plan type controls			X	
N (plan-block-years, deadpooling only)	111,097	111,097	111,097	
Unique blocks	12,478	12,478	12,478	
Subsidiaries	291	291	291	
Parent companies	207	207	207	

Notes: Identical to Table 2 in sample, weighting, covariates, estimator (`dripw`) and base period (`long2`); point estimates are therefore the same. Standard errors in parentheses are clustered at the parent company, with each block assigned the parent group under which it carried the most covered lives over its life. "Parent companies" counts every parent group that appears in the arm's sample in any year, which exceeds the number of modal-parent clusters. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.6 How the two estimators differ

Both are doubly-robust difference-in-differences estimators (Sant'Anna and Zhao 2020): each combines an outcome-evolution model with a propensity-score model for the probability of treatment, and is consistent if either one is correct. They differ only in how that propensity score is estimated.

`dripw` fits the propensity score by ordinary logit maximum likelihood, choosing $\hat{\gamma}$ to maximise the likelihood of the observed treatment assignment, and then forms inverse-probability weights from the fitted $\hat{p}(X) = \Lambda(X'\hat{\gamma})$. Balance across the covariates is a by-product: nothing in the objective requires it.

`dripw` instead uses *inverse probability tilting* (Graham, Campos De Xavier Pinto, and Egel 2012), which chooses $\hat{\gamma}$ to solve the moment condition

$$\sum_i (1 - D_i) \exp(X_i' \hat{\gamma}) X_i = \sum_i D_i X_i, \quad (16)$$

so the reweighted control group matches the treated group's covariate means *exactly*, by construction. Balance is the objective rather than a consequence.

The practical difference appears when a covariate nearly separates treatment from control. Consider the plan-letter dummies: Plans M and N did not exist before June 2010, so for a cohort whose comparison window sits entirely in 2007–2010 those dummies are identically zero among the treated. Logit responds to perfect separation by sending the coefficient to $\pm\infty$; the fitted propensity approaches 0 or 1, the inverse-probability weights diverge, and the cell is dropped. The tilting condition (16) has no such failure mode, because it asks only that reweighted means agree, which remains solvable. The same mechanism explains why `dripw`'s standard errors are several times larger in every covariate-adjusted specification while

being *identical* to `dripw`'s when no covariates are included—with no propensity model to fit, the two estimators coincide exactly.

C.7 What this costs in sample

Because whole cohort-period cells are dropped rather than down-weighted, the loss is concentrated in the early cohorts: under the fully-adjusted specification, `dripw` discards every cell whose comparison window touches 2007–2010. In the fully adjusted premium specification, `dripw` omits 304 of 760 group-time cells on the non-deadpooling arm and 280 of 722 on the deadpooling arm. `dripw` omits 68 and 2.

We therefore do not adopt `dripw` as the main specification: it answers a question about a later, smaller panel than the one we observe. Reported alongside, it leaves the substance unchanged. On the deadpooling arm of Table 2, `dripw` returns a premium effect of \$1,362 against `dripw`'s \$1,005, and a profit effect of \$548 against \$398—larger under the default estimator, as the weight instability above would predict, but the same sign and the same order of magnitude.

The pre-trend diagnostics deserve a word of caution, because they run the other way and it is tempting to read them backwards. Under `dripw` the premium pre-trend test returns $p = 0.067$, against 0.004 under `dripw`: the default estimator appears to pass where our preferred one fails. That is not cleaner identification. The same inflated variance that widens `dripw`'s confidence intervals also widens the covariance the Wald statistic divides by, so the test loses power and fails to reject a violation that is present either way. A diagnostic that improves because the estimator became noisier is not evidence in the estimator's favour. The loss ratio survives under both (0.260 and 0.619), which is the outcome we lean on regardless.

The estimator choice therefore moves precision and sample, not conclusions.

C.8 The Two-Way Fixed-Effects Analogue

This appendix re-estimates the closure effect with the conventional two-way fixed-effects analogue, which the main text does not use, and reports whether it agrees. It is run with year, state, plan-type and plan fixed effects added in turn. It is not an independent test of identification (with staggered timing, TWFE weights some 2×2 comparisons negatively (Goodman-Bacon 2021)) so it is reported only to show that the result is not an artifact of the Callaway–Sant'Anna machinery. It agrees with the main-text magnitudes and is stable as the fixed effects are added.

Table A7 reports the before-and-after-closure comparison for premiums, claims, and profits, with never-closed plans as the control group, progressively adding year, state, plan-type, and plan fixed effects across columns. The estimates agree with the main-text magnitudes.

C.9 The Quoted-Price Effect Under Alternative Age Weightings

The quoted-price index in Figure 5 weights ages 65 to 85 by the MCBS age distribution of Medigap holders, so it answers what closure does to the price faced by the age mix that actually holds these policies. That weighting is a choice, and Figure A3 shows what the alternatives would give: the price at a single age, 65 or 75, and a flat average over the whole 65–85 range. The four series are the same estimation on the same sample and differ only in how the age cells are combined, so any difference between them is the weighting and


Table A7. Premiums and Profits Rise After Closure in Every Fixed-Effects Specification

	(1)	(2)	(3)	(4)	(5)	Mean before closure
<i>Panel A: Per-Person Premiums (2025 \$)</i>						
NAIC	\$764.37*** (96.76)	\$799.01*** (111.41)	\$898.41*** (109.75)	\$590.03*** (120.87)	\$482.55*** (119.97)	\$2,615
NAIC, CSG sample	\$737.51*** (110.28)	\$738.71*** (113.01)	\$886.65*** (112.98)	\$446.37*** (126.32)	\$366.08*** (55.47)	\$2,405
CSG	\$299.47*** (94.83)	\$301.00*** (100.77)	\$452.92*** (77.46)	\$100.33 (86.22)	\$169.88*** (60.63)	\$2,976
<i>Panel B: Per-Person Claims (2025 \$)</i>						
NAIC	\$483.42*** (69.68)	\$507.95*** (77.33)	\$582.35*** (73.19)	\$361.52*** (83.15)	\$347.90*** (84.19)	\$2,035
NAIC, CSG sample	\$488.38*** (84.71)	\$487.60*** (85.01)	\$583.75*** (82.69)	\$273.85*** (95.62)	\$324.56*** (46.76)	\$1,884
<i>Panel C: Per-Person Profit (2025 \$)</i>						
NAIC	\$280.96*** (39.97)	\$291.06*** (43.64)	\$316.06*** (39.82)	\$228.50*** (40.83)	\$134.66*** (41.95)	\$580
NAIC, CSG sample	\$249.13*** (38.76)	\$251.11*** (39.81)	\$302.90*** (35.30)	\$172.52*** (36.90)	\$41.52 (29.23)	\$521
<i>Panel D: Loss Ratio (percentage points)</i>						
NAIC	-3.62*** (1.04)	-3.57*** (0.89)	-3.59*** (0.59)	-2.80*** (0.59)	0.03 (1.05)	79.0
NAIC, CSG sample	-3.05*** (0.83)	-3.13*** (0.78)	-4.05*** (0.63)	-2.48*** (0.71)	2.04* (1.16)	79.1
Year FE		X	X	X	X	
State FE			X	X	X	
Plan Type FE				X	X	
Plan ID FE					X	
N (Life-Years)	174,112,756	174,112,756	174,112,756	174,112,756	174,112,754	
N (Plan-Years)	214,895	214,895	214,895	214,895	214,893	

Notes: All regressions are two-way fixed-effects specifications weighted by total lives, with never-closed plans as the control group and standard errors clustered at the NAIC insurer level. All monetary units are per enrollee in 2025 dollars. Standard errors are in parentheses, and stars denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The result is that premiums and profits rise after closure and the estimates are stable as year, state, plan-type, and plan fixed effects are added, so the increase is not an artifact of which plans close, where, or when. See the text above for why this is a robustness exhibit rather than the paper's preferred estimator.

nothing else.

Figure A3. The Quoted-Price Closure Effect Under Four Age Weightings



Notes: Callaway–Sant’Anna ATT of plan closure by event time on the CSG quoted price, normalized to zero at $t = -1$, on the specification of Figure 5. Navy is the MCBS-weighted index reported there. The other three series reweight the same age cells: the quoted price at age 65 alone, at age 75 alone, and a flat unweighted average over ages 65 to 85. All are annual rates in January 2025 dollars for a non-tobacco applicant, and all are the female quote, which is the gender the CSG age profile is built on. Confidence intervals are drawn for the MCBS series only, and the single-age series are offset horizontally by a tenth of a year, to keep the figure readable. Joint pre-trend tests (leads $t - 6$ to $t - 2$, $t - 1$ the base period): no pre-trend test in this figure is significant at the 5% level (the MCBS-weighted index, $p = 0.185$; the equal-weight index, $p = 0.152$).

C.10 Netting Out Aging from the Stayer Burden

A simple measure reads a trapped stayer’s burden off the growth of a closed block’s per-enrollee premium. This appendix nets out the part of that growth that is the pool aging.

The aging confound, and how far it goes. A closed block issues no new policies, so its mean attained age rises by about a year every year; an open block refreshes with new sixty-five-year-olds and holds its mean age roughly flat. Under attained-age rating the closed block’s *pool-average* premium therefore rises even if the insurer never raises a single rate, and the block-level wedge $g_{dp} - g_{cf}$ mixes this mechanical aging drift together with adverse retention and with genuine rate increases. The aging drift, crucially, is *not* deadpooling-specific harm: a trapped stayer would have aged along the same rate schedule in an open block too, so the excess a stayer pays *because* the block deadpooled should net it out. We bound how much of the wedge it accounts for.

The design exploits rating type. Community-rated and issue-age plans are immune to the individual aging gradient by construction (a community-rated premium does not vary with age, and an issue-age premium is fixed at the age of purchase) whereas an attained-age premium climbs as the enrollee ages. Rating type is not randomly assigned, however: open-market states are almost entirely community-rated, and stronger switching protection correlates with community rating, so a raw comparison across rating types would confound aging with the regulatory regime. We therefore hold the regime fixed and run the entire comparison *within* the no-switching (maximal lock-in) regime, which is the relevant regime for a trapped stayer and which contains all three rating types. Within it, the raw attained-age wedge is 4.7 percentage points a year. Netting out only

the individual age gradient—which we read directly off the CSG quoted-rate curve as 2.8 points per year of age, and which is *identical* for deadpoolers and non-deadpoolers (2.8 points, insignificant net of state, plan, and year fixed effects, with no steepening after closure)—leaves a deadpooling-specific wedge of 2.0 points. Because the attained-age curve is no steeper in deadpooling blocks, a stayer rides the same age schedule whether or not the block closes, and netting it is legitimate. Benchmarking instead against the aging-immune blocks leaves 1.5 points, but this over-corrects: issue-age blocks, which dominate the immune group, cannot reprice their aging pool, so as their pool sickens their markup mechanically falls and their wedge understates the harm an attained-age stayer faces. Table A8 carries the three treatments through to the fifteen-year stayer present value.

Table A8. The Paper’s Headline Stayer Burden: \$5,735 Over Fifteen Years After Netting Out Aging

Treatment of aging		Wedge (pp/yr)	Stayer PV, 15 yr
None (uncorrected)	upper bound	4.7	\$12,216
Individual age gradient netted	headline	2.0	\$5,735
Benchmarked to aging-immune blocks	lower bound	1.5	\$2,937

Notes: All three rows are estimated within the no-switching regime, where the raw fifteen-year stayer present value (\$12,216) is somewhat below the \$12,980 full-sample figure. “Individual age gradient netted” subtracts 2.8 percentage points per year (the CSG quoted-rate 65-to-85 gradient in attained-age plans, which is the same for deadpoolers and non-deadpoolers) from the counterfactual, leaving the deadpooling-specific rent; it is our preferred correction. “Aging-immune” benchmarks the deadpooled ramp against community- and issue-age-rated blocks, and is a lower bound because issue-age blocks cannot reprice an aging pool. The present value discounts at 3% and weights each year by survival at the 3.9% annual mortality hazard. The wedge and the present value need not order identically across rows: the aging-immune benchmark removes more than the individual age gradient does, so its wedge can exceed the headline wedge while its present value remains the smaller of the two, because the two corrections also differ in the base premium the wedge compounds on. Mechanical aging accounts for roughly two-thirds of the headline stayer figure—a large correction that nonetheless leaves a substantial deadpooling-specific residual rather than erasing it. The netting is legitimate precisely because the quoted attained-age gradient is identical for deadpoolers and non-deadpoolers: a stayer rides the same age schedule whether or not their block closes, so deadpoolers extract more through faster across-the-board rate increases rather than through a steeper age curve.

The preferred correction cuts the burden to roughly a third, from \$12,216 to \$5,735 over fifteen years, with the aging-immune benchmark placing a floor at \$2,937. Mechanical aging thus accounts for something like two-thirds of the headline stayer figure—a large correction, but one that leaves a substantial deadpooling-specific residual rather than erasing it. The one apparent sign that deadpoolers milk attained-age blocks *harder* (the aging-invariant markup ramp is about -2.9 points steeper in attained-age than in aging-immune deadpooled blocks) does not survive scrutiny: it reflects issue-age blocks’ inability to reprice their aging pool, not a steeper age curve, since the quoted attained-age gradient is identical for deadpoolers and non-deadpoolers. Deadpoolers do extract more rent from attained-age blocks, but through faster across-the-board rate increases, which the wedge already captures, not through a steeper age schedule, which it would wrongly remove.

Aging makes the premium-ramp estimates the least secure numbers in the paper. We therefore quote the aging-netted \$5,735 as the point estimate and carry the interval from \$2,937 to \$12,216 with it, not presenting any of the three as precise.

The transfer estimates built on the loss ratio do not share this weakness, but they answer a narrower question. A loss ratio is claims over premiums, and an older or sicker pool raises both, so the ratio is invariant to the pool’s composition, including its age. That invariance is what makes the aggregate transfer credible where the per-enrollee ramp is fragile, and that the mechanical rebate reproduces the causal aggregate to within a few percent is strong evidence the transfer is real. But invariance to pool composition cuts both ways: a measure that cannot see the pool deteriorate cannot price the cost of being sorted into a deteriorating pool. The loss-ratio route recovers the extraction channel and, by construction, sets the segmentation channel to zero. It is

thus a lower bound on the burden a trapped stayer bears, not a competing estimate of it—which is why we retain the premium ramp, aging correction and all, as the measure of individual harm, and rest the *aggregate* claim on the loss ratio.


C.11 Heterogeneity by Market Share: Large-Company Deadpoolers versus Never-Closing Plans

The closure event study above pools every closing plan. Because the strategic close-and-reopen pattern is concentrated among the large national parents, we sharpen the treated group to the four large companies we identify as most commonly engaging in this practice, and re-estimate the same Callaway–Sant’Anna design against *all never-closing plans* as the control group, adjusting for state and plan-type through doubly-robust regression. Figure A4 shows the event study for profit, the loss ratio, premiums and claims, and enrollment. The pattern matches the pooled estimates but is larger: per-person premiums climb steadily after closure, reaching \$820 per enrollee three years out and \$1,149 by five years, while claims rise by less (\$524 at three years). The residual accrues to profit, which runs *below* the never-closer benchmark throughout the pre-closure window (–\$165 five years out), crosses it around the closure year, and climbs to \$499 per enrollee by five years; the loss ratio falls by 0.054 at three years and 0.088 at five. Enrollment declines steadily and substantially over the same window: five years after closure the large parents’ closed blocks retain 8% of their year-before-closure enrollment, against 44% for other deadpoolers (panel D).

One caveat tempers the magnitudes. Unlike the pooled closure event study, the pre-closure trends here are not flat. Measured against a universal base period, all 4 estimated pre-periods have confidence intervals excluding zero for premiums, claims and profit alike: each of the three sits well below the never-closer path five years out (–\$462 for premiums, –\$297 for claims, –\$165 for profit) and rises steadily toward it as closure approaches. The loss ratio and enrollment are the exceptions, with no pre-period distinguishable from zero. Because the post-closure levels are normalized against a pre-period that is itself trending, we read these sharpened estimates as corroborating the pooled closure results (Figure 3) in direction and mechanism rather than as a tighter point estimate of them. As an even sharper single-firm benchmark, Appendix C.11.1 compares the major deadpoolers against a single benchmark carrier (the market’s largest continuously-open, never-deadpooling operator) and finds the same profit and markup wedge.


Is the effect a large-company phenomenon? The comparison above sharpens the treated group to four large parents, which raises the obvious question of whether deadpooling pays only at that scale. Figure A4 answers it by estimating the same event study a second time on *every other* deadpooler (same never-closing control group, same covariates, same estimator) and overlaying the two. Both groups reprice after closure, so the mechanism is not confined to the largest firms; but the large parents extract roughly twice as much. Five years out they raise premiums by \$1,149 per enrollee against \$581 for the rest, and profit by \$499 against \$249. The loss-ratio gap is starker still: -0.088 against -0.011, meaning the smaller deadpoolers’ premium increases roughly track their claims while the large parents’ do not. Scale is therefore not what makes deadpooling possible, but it is strongly associated with how hard the closed block is worked.

Figure A4. Both Large and Small Deadpoolers Reprice After Closure, but Large Parents Extract About Twice as Much



Notes: Two Callaway–Sant’Anna event studies drawn on one set of axes. The solid navy series is closures by the four large parents most associated with deadpooling; the dashed maroon series is closures by all *other* deadpooling blocks. Both use the identical control group (all never-closing plans), covariates (state and plan type, doubly-robust), estimator (`drimp` with a universal base period), clustering (NAIC) and weighting (lives), so the two lines differ only in who is treated. Confidence intervals are offset horizontally so they do not overlap. Panels are per-enrollee profit (A), the loss ratio (B), per-enrollee premiums and claims (C; premiums dark, claims light) and enrollment (D). Panel D’s outcome is a block’s enrollment as a percent of its enrollment in the year before closure (never-closing blocks: of their mean enrollment), weighted by that base-year enrollment, and is plotted as 100 plus the ATT; blocks that run out of enrollees leave the panel, so it is measured among blocks still reporting. The pre-period is not flat for the large-parent series, which fails the joint pre-trend test for premiums, claims, profit and the loss ratio; the smaller deadpoolers’ series passes it for premiums, claims and profit and fails it only for the loss ratio, so the contrast between the two post-closure paths is the more reliable comparison here than either level on its own. Joint pre-trend tests (leads $t - 6$ to $t - 2$, $t - 1$ the base period): tests reject at the 5% level for premiums, large parents, $p = 0.000$; claims, large parents, $p = 0.000$; profit, large parents, $p = 0.000$; the loss ratio, large parents, $p = 0.023$; the loss ratio, other deadpoolers, $p = 0.024$; all other pre-trend tests in this figure are not significant at the 5% level (premiums, other deadpoolers, $p = 0.222$; claims, other deadpoolers, $p = 0.516$; profit, other deadpoolers, $p = 0.095$; enrollment, large parents, $p = 0.225$; enrollment, other deadpoolers, $p = 0.505$).

Figure A5. The CSG Quoted Price Rises After Closure for Large and Small Deadpoolers



Notes: The design of Figure A4 run on the CSG quoted-price index (a fixed basket over ages 65–85 weighted by the MCBS age distribution of Medigap holders, annual 2025 dollars), on the CSG sample of Table 2: block-years with a CSG quote, with never-closing controls only where the block is quoted in its own right. Solid navy is the four large parents, dashed maroon all other deadpoolers; same controls, estimator, clustering and lives weights as Figure A4. Because the basket holds the age mix fixed, the rise cannot come from an aging pool. Five years after closure the index is \$890 higher for the large parents and \$619 for other deadpoolers. The large-parent series fails the joint pre-trend test (leads $t - 6$ to $t - 2$, $p = 0.000$): its quoted price is already rising toward the control path before closure, so its post-closure level partly continues that trend. The other deadpoolers' series passes it ($p = 0.251$).

C.11.1 A Single-Firm Never-Deadpools Benchmark

The comparisons above use all never-closing plans as the control group. A concern with any broad control is that it pools plans of very different insurers, and the composition of that pool can itself vary by region. It also raises a comparability question: if small carriers lack the resources to stand up subsidiaries and deadpool at all, then a control group drawn from all never-closing plans is not the right counterfactual for a firm that does. As a sharper single-firm benchmark we therefore compare deadpoolers against one large national carrier (one with every resource needed to deadpool, which nonetheless does not) whom we do not name. It runs a single continuously-open block in each market, so it absorbs the same calendar-time cost shocks and the same ordinary within-block aging without the strategic close-and-reopen segmentation of healthy versus sick. We re-estimate the same Callaway-Sant’Anna design with that carrier as the *sole* never-treated control, on two profitability outcomes — the loss ratio (claims/premium) and per-enrollee dollar profit. Two treated definitions give nearly identical answers: all deadpoolers, and only the four large companies we identify as most commonly engaging in this practice. Deadpooled blocks run a 4-percentage-point lower loss ratio and earn roughly \$410–\$460 more per enrollee than the benchmark’s open blocks (Table A9). The loss-ratio series passes the joint pre-trend test ($p = 0.484$), but the profit series does not ($p = 0.000$): deadpoolers begin *below* the benchmark in per-enrollee profit and are already catching up before closure, so the profit comparison should be read with that pre-trend in mind.

Table A9. Deadpooled Blocks Out-Earn the Benchmark Carrier’s Open Blocks by \$410–460 per Enrollee

Outcome	Overall ATT vs. the benchmark carrier	
	Major-parent deadpoolers	All deadpoolers
Loss ratio (claims/premium)	-0.040 (0.027)	-0.041* (0.024)
Per-enrollee profit (\$)	457*** (151)	407*** (128)

Notes: Callaway-Sant’Anna (2021) overall ATT of deadpooling on each profitability wedge, with the benchmark carrier (a never-deadpools operator) as the sole never-treated control. The first column restricts to the four large companies we identify as most commonly engaging in this practice; the second is all deadpoolers. Loss ratio is claims/premium, per-enrollee profit is in dollars; all outcomes winsorized at 1/99. Estimates are unconditional — the doubly-robust state and plan-type adjustment is not identified against the thin single-firm control (it drops the non-overlapping pre-period cells and inflates the post ATT). Because no covariates enter, *csdid*’s two doubly-robust estimators coincide exactly here and the distinction between *drip* and *dripw* does not arise; the table is run under the default (*dripw*). Standard errors (parentheses) are NAIC-clustered. All estimates are enrollment-weighted (Callaway-Sant’Anna importance weights set to total lives). Against the sharpest available benchmark (the market’s largest operator, which runs one continuously-open block per market and never deadpools) deadpooled blocks run a 4-point lower loss ratio and earn \$410–460 more per enrollee. The answer is nearly identical whether treatment is all deadpoolers or only the four major parents, so the result is not an artifact of the big four. Stars denote significance of the overall ATT: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Joint pre-trend tests (leads $t - 6$ to $t - 2$, $t - 1$ the base period): the test rejects at the 5% level for large-parent profit ($p = 0.000$); all other pre-trend tests in this table are not significant at the 5% level (large parents: loss ratio $p = 0.484$; all deadpoolers: loss ratio $p = 0.579$, profit $p = 0.159$).

C.12 Letter-Matched Controls

Appendix C.5 identifies plan letter as the covariate that breaks *csdid*’s default propensity model. There are two ways to hold plan letter fixed. One is to adjust for it, which is what the covariate specification does and what *drip* makes feasible. The other is to match on it, comparing each closing block only to never-closing blocks of the same letter, so that the letter never enters the propensity model at all. The two routes share no estimating machinery, which makes their agreement (or disagreement) informative about the adjustment rather than about either estimator. This appendix reports the matched version, as a robustness check on the covariate-adjusted closure estimates of Table 2 in the main text.

C.12.1 The interaction is not available as a covariate

The natural way to let each plan letter carry its own time path is a letter $\times$ year interaction. That is not estimable inside `csdid`: the documentation is explicit that only base-period covariate values are used, and within a single $ATT(g, t)$ the base period is a single year, so a letter $\times$ year term collapses to letter and adds columns collinear with the letter dummies already present.

The saturated cross-section version, `i.state##i.standnum`, is estimable in principle and fails in practice. Under `drimp` it separates completely: 100% of group-time cells are omitted in all 6 fits (premiums, profit and loss ratio, both panels) and the command reports a sample size of zero.

The reason is arithmetic rather than numerical, and it follows from where `csdid` fits the propensity model: separately inside each $ATT(g, t)$ cell, using cohort g 's treated blocks against the never-treated pool. The saturated interaction contributes one column per occupied state $\times$ letter stratum, of which there are 472. The never-treated pool is not the binding constraint—it occupies 465 of those 472 strata, so controls are available almost everywhere. The treated side is. Of the 40 closure cohorts, 31 have *fewer treated blocks than there are strata* (median 190 blocks against 472 columns), so for those cells the propensity model has more parameters than treated observations and nothing is identified. The larger cohorts fail for the second reason: a cohort's treated blocks occupy a median of only 27.2% of the strata, and at most 71.8%, so the remaining strata contain never-treated blocks and no treated one. Each such dummy predicts non-treatment perfectly, which is complete separation, and the tilting problem has no interior solution. Both counts are computed pooling over time, so they *overstate* the overlap available inside any single (g, t) cell, where the treated blocks must also be observed in both the base period and t . Inverse probability tilting extends how much conditioning a propensity model will bear, but it cannot manufacture treated variation in a stratum that contains none. The interaction therefore has to come from the design.

C.12.2 Matching the control group on plan letter

We estimate the closure effect separately within each standardized letter, using `drimp` and never-closing blocks of that same letter as controls, and combine the ten estimates with enrollment weights. Because treated and control blocks share the letter, any letter-specific year effect differences out exactly. Table A10 reports the result, with and without state fixed effects inside each stratum, against the Column 3 estimates of Table 2 that it is checking.

For deadpooling closures the matched premium estimate is \$1,012, against \$1,005 from the fully adjusted specification of the main-text closure estimates (Table 2)—a difference of \$7, or 0.7%. Profit agrees to \$4 (0.9%). Two procedures with no machinery in common, one conditioning on plan letter and one holding it fixed by construction, land within one percent of each other. The exception is the non-deadpooling premium, where the matched estimate of \$365 sits 10.8% above the adjusted \$329; non-deadpooling profit agrees to 1.6%. The deadpooling–non-deadpooling contrast that the paper rests on is, if anything, slightly larger under matching.

Because the ten strata partition the blocks and are estimated separately, the pooled standard error must account for correlation across strata induced by insurers operating in several of them at once; we obtain it by clustering the stacked influence functions on insurer rather than re-estimating.

We do it that way because a literal block bootstrap would require roughly sixty `csdid` fits per replication

Table A10. Letter-Matched Controls Reproduce the Covariate-Adjusted Estimate

Outcome	Non-deadpooling		Deadpooling	
	Letter-matched	Letter-matched	Letter-matched	Letter-matched
Per-person premiums	365 (145)	333 (146)	1,012 (199)	954 (202)
Per-person claims	185 (82)	147 (88)	610 (114)	579 (126)
Per-person profit	180 (68)	186 (66)	402 (89)	375 (79)
Loss ratio	-0.025 (0.010)	-0.029 (0.011)	-0.057 (0.017)	-0.054 (0.015)
State FE		X		X

Notes: Each entry is an enrollment-weighted combination of ten within-letter Callaway–Sant’Anna estimates, each comparing closing blocks of one standardized letter to never-closing blocks of the same letter, estimated by `drip`. The second column of each pair adds state fixed effects inside the stratum, and is reported as a check rather than as an improvement: within a single letter `i.state` is 47 dummies fitted on roughly a tenth of the panel, so the doubly-robust estimator drops cohort-period cells for want of overlap and the state-adjusted estimate rests on a smaller base than the letter-matched one beside it. Where the two columns disagree, the difference can therefore reflect a changed sample as much as a changed adjustment. Parenthesised figures are standard errors clustered on the insurer, computed from `csdid`’s influence functions rather than by re-estimation: each block’s influence on its letter’s aggregate is scaled by that letter’s enrollment weight and summed within insurer, so an insurer’s contributions from different letters add before being squared. That cross-letter covariance is what makes these standard errors correct rather than conservative—the letter strata partition the blocks, but insurers sell many letters, so treating the strata as independent understates the variance. The procedure was validated by reproducing `csdid`’s own within-letter estimates and standard errors exactly. A 999-replication Mammen multiplier-bootstrap percentile interval over insurers was also computed and agrees with these standard errors; it is not shown, to keep one measure of uncertainty per cell. Clustering raises the standard errors by factors of 1.02 to 1.48, and by more in the deadpooling panel (up to 1.48) than the non-deadpooling one (up to 1.11), which is expected: deadpooling is concentrated in a few large parents that sell across many letters. 15 of the 16 intervals exclude zero; the exception is per-person claims for non-deadpooling closers with state fixed effects, which is the one cell in the table that cannot be distinguished from zero. Plan M is retained for completeness though its estimates are unreliable; excluding it changes no entry in this table at the precision shown.

(several hundred hours) while the influence functions `csdid` already computes deliver the same clustering in one pass.

Within a single letter `i.state` is 47 dummies fitted on a tenth of the panel, so the doubly-robust estimator drops cohort-period cells for want of overlap and the surviving estimate rests on a smaller base than the letter-matched specification. That is why we report the letter-matched estimate as the primary one and treat the state-adjusted variant as a check rather than an improvement.

C.12.3 Matching does not repair the pre-trends

The matched design leaves the parallel-trends verdict of Appendix C.1 intact. In the deadpooling panel the premium pre-trend fails for Plan F, which carries the most enrollment of any letter in the paper, under both specifications and at $p < 0.001$; only 2 of the ten letters pass, and adding state fixed effects changes no verdict there. The non-deadpooling panel is less stable—4 letters pass under letter matching alone and 3 once state effects are added, with three letters changing side, Plan F among them in the favourable direction. That last flip is worth isolating: Plan F loses no cells to the state dummies in either panel, so its non-deadpooling reversal is an adjustment effect rather than a changed sample. It has no counterpart in the deadpooling panel, which is the one the paper’s claim rests on. Whatever generates the pre-closure premium divergence there, it is not a letter composition effect that matching removes, and it is not geography.

Plan G is the one large letter whose premium path is flat before closure, and comfortably so. G is also the letter with the least statutory interference in its availability over the sample period, which is consistent with

the pre-trend being bound up with the changing menu of plans rather than with closure itself. We note the pattern without resting anything on it: one letter is one letter.

The loss ratio again behaves better than the levels, though less clearly than on the full panel. It satisfies parallel trends for Plan F under both specifications, and in the non-deadpooling panel for 5 of the ten letters under letter matching alone and 4 with state effects added. Its one reversal is instructive: Plan G's loss-ratio pre-trend passes in the letter-only specification and fails once state effects are added—in exactly the letter that loses a quarter of its cells to those effects. That is the failure mode of Appendix C.5 reappearing at letter scale, and we read the reversal as a consequence of the deleted cells rather than as evidence about Plan G.

D Appendix: How Parent Companies Carry Out Deadpooling

D.1 Closures Are Coordinated, Subsidiary-Wide Events

The main text treats deadpooling as a deliberate, coordinated action rather than the incidental closure of a single plan. If that reading is correct, a subsidiary's closures should cluster: when a subsidiary shutter one plan it should be shuttering its other plans at the same time, both across standardized letters within a state and across states. We test this directly with a discrete-time closure hazard at the block level, where a block is a subsidiary $\times$ state $\times$ standardized letter $\times$ (individual/group) cell (136,627 block-years at risk, 11,249 first-closure events). For each block-year we compute two *leave-one-out* sibling closure rates (the share of the *same* subsidiary's *other* letters in the *same* state that close that year, and the share of the same subsidiary's blocks in *other* states that close that year) and ask whether either predicts the block's own closure. Because a market-wide regulatory shock (the 2020 prohibition on selling Plans F/C to newly-eligible beneficiaries, or the 2010 introduction of Plans M/N alongside the ACA) closes a given letter for every insurer at once, we absorb letter $\times$ year and state $\times$ year fixed effects, so the sibling coefficients are identified off *within*-subsidiary co-movement beyond the market. Standard errors are clustered by subsidiary.

Closures are overwhelmingly coordinated on both margins. Across letters, when a subsidiary closes a plan in a state, on average 73.7% of its other letters in that state close the same year, against 2.3% in years the block stays open; across states, 61.3% of a subsidiary's other-state blocks close in the same year as a given closure, against 3.5% otherwise. Both gaps survive letter $\times$ year and state $\times$ year fixed effects and stay large when the two channels are entered jointly (52.6 and 45.8 percentage points), and restricting the signal to *deadpooling* closures gives the same picture. Closure is thus a subsidiary-wide event: when a parent winds a subsidiary down it freezes the entire book (every letter, in every state) in a single year.

The cross-state pattern carries a further implication for *why* blocks close, and one that matters for the event-study designs in the main text. Because a subsidiary's blocks close together across states (and those states differ in their local competitors, demographics, regulatory regimes, and realized claims experience) the timing of any one block's closure cannot plausibly be a reaction to that block's own pool. A closure driven by a particular pool's deteriorating experience would arrive idiosyncratically, whenever and wherever that pool happened to sour; the observed nationwide simultaneity instead points to a firm-level decision taken on a common corporate clock, exogenous to any individual block's realized costs or health mix. In other words, the subsidiary is not closing a plan in response to that plan's pool, but shutting its whole book at once for reasons that operate above the level of the individual block. This is not to say closure is random or non-strategic (the

whole thrust of the paper is that deadpooling is a deliberate firm-level strategy) but that the strategic decision is taken at the subsidiary level and is exogenous to any *individual* block’s realized experience; the two statements are complements, not tensions. This supports treating the timing of closure as exogenous to a block’s own pool (the identifying assumption underlying the closure event studies) rather than as an endogenous response to the very experience whose post-closure change we measure.

D.2 The Timing of Closure and Parent Reentry

Section 4.1 states the two headline exposure shares: 17.6% of life-years through 2020 are spent in a plan that deadpools, and about 38% of beneficiaries in a closing plan see their parent reenter within five years. By the closure year itself, 26% of these enrollees have seen their parent reenter, so the closure year accounts for the largest single share of the five-year figure. This appendix adds two details of that timing, which describe the construction rather than a result.

The timing also shows that reentry often *precedes* closure. 12% of enrollees in a closing plan are already in a block whose parent has opened a successor in the year before the closure takes effect—close to half of the 26% observed by the closure year itself. The successor is therefore frequently in place before the predecessor shuts, which is what one would expect if the closing and the opening are planned together rather than the opening being a response to the closing.

Reentry windows. Our preferred definition counts a parent as reentering if it opens a block in the same market within five years of a closure, matching the five-year reentry ban. Table A11 recomputes the exposure shares with one-, three- and seven-year windows. The deadpooling share is not driven by reentries late in the five-year window: shortening it to three years lowers the share of life-years through 2020 spent in plans that deadpool from 17.6% to 16.2%, and shortening it to one year (from the year before closure through the year after) lowers it to 13.3%. About three quarters of the headline share therefore comes from parents that reopen within a year of closing. Lengthening the window to seven years raises the share to 21.0%, so the five-year definition does not count reentries that come just after the ban expires.

Table A11. Deadpooling Exposure by Reentry Window


Reentry window	Life-years in plans that deadpool (%)		Life-years in a closed deadpooled block (%)	Enrollees in closing plans whose parent reenters (%)
	Through 2020 (1)	Fully observed (2)		
1 year	13.3	12.9 (2024)	7.1	28.6
3 years	16.2	16.0 (2022)	8.7	34.6
5 years (preferred)	17.6	17.6 (2020)	9.4	37.6
7 years	21.0	21.3 (2018)	11.1	44.3

Notes: A block deadpools if it closes and its parent opens a block in the same state, plan letter and group-or-individual market between the year before closure and N years after it. All shares are weighted by covered lives. Column (1) restricts both numerator and denominator to years through 2020, as in the headline share. Column (2) extends that restriction to the last year in which an N -year window is fully observed (in parentheses). Column (3) is the share of all life-years in the panel spent in a deadpooled block after it closes. Column (4) is the share of enrollees in plans that close through 2020 whose parent reenters within the window; for the seven-year window it is restricted to closures through 2018, so that the window is fully observed. The five-year row reproduces the shares reported in Section 4.1 and the introduction. Source: `analysis/deadpooling_window_robustness.do`.

Figure A6 shows how quickly plans close in the first place, tracing the share of enrollees whose plan is still open n years out. Closure is the common case rather than the exception, and the decline is steepest in the first


four years.

Figure A6. Nearly Half of Enrollees Are in a Plan That Closes Within Ten Years



Notes: Enrollment-weighted (by total lives): each point is the lives-weighted share of enrollees whose plan is still open $[n]$ years in the future. Restricted to plans first appearing before 2015, so each can be followed a full ten years. The two slopes annotated on the figure are descriptive OLS slopes of the plotted series over event years 0–4 and 5–10 respectively; they are not estimates of a treatment effect. Closure is the common case, not the exception: 18% of enrollees are already in a closed plan at the left edge, and only 53% remain in a plan still accepting new members ten years out. The decline is steepest in the first four years, at about five percentage points a year, flattening to roughly one point a year thereafter.

Figure A7. Agent Knowledge of the MA-to-Medigap Underwriting “One-Way Door”



Notes: Share answering correctly that a client who switches from Medicare Advantage to Medigap later can be medically underwritten (and denied or surcharged). Agents who pass the survey’s attention and invisible-text checks; $N = 306$ agents answered this item. Question asked: “A healthy 67-year-old in a Medicare Advantage plan asks: ‘If I try MA now, can I switch to a Medigap plan in a few years if I don’t like it?’ How would you answer this question?”

D.3 Brand Distance Between Parent and Subsidiary

The mechanism requires that enrollees not connect the block that closed to the fresh block their parent opens next door. If that opacity were deliberate, it should leave a firm-side trace: a parent would route the strategy through subsidiaries whose names do not resemble its own. It does not. We classify each subsidiary as brand-similar or brand-dissimilar to its parent by stripping generic tokens from the parent’s legal name in the NAIC parent–subsidiary map and testing whether any distinctive token survives in the subsidiary’s name; ties and wholly generic parent names are coded *similar*, so the dissimilar group is never inflated by ambiguous cases.

The raw gap is composition. Dissimilar-named subsidiaries deadpool 31.0% of their plans against 25.3% for similar-named ones. That gap is between parents, not within them: Table A12 absorbs parent, state and plan-letter effects and the difference disappears. Deadpooling is a strict subset of closure, so the two rows of Panel A carrying opposite signs is itself a sign that neither is a real effect. Panel B asks whether a dissimilar-named block reprices its trapped pool harder, and finds nothing. Restricting to block-years after an acquisition (so that brand distance reflects a decision the current owner actually made) does not change the answer.

The two panels of Table A12 estimate, for plan p of subsidiary $s(p)$ under parent $f(p)$, in state $k(p)$ and plan letter $\ell(p)$,

$$Y_p = \beta D_{s(p)} + \alpha_{f(p)} + \gamma_{k(p)} + \delta_{\ell(p)} + \varepsilon_p, \quad (17)$$

$$P_{pt} = \beta (D_{s(p)} \times Post_{pt}) + \beta_0 D_{s(p)} + \beta_1 Post_{pt} + \lambda C_{pt} + \alpha_{f(p)} + \gamma_{k(p)} + \delta_{\ell(p)} + \tau_t + \nu_{pt}, \quad (18)$$

where D_s indicates a subsidiary whose name shares no brand token with its parent’s, Y_p is an indicator that the plan deadpools (or ever closes), P_{pt} is per-enrollee premium, $Post_{pt}$ indicates years at or after the plan’s closure, and C_{pt} is per-enrollee claims. Both are weighted by covered lives with standard errors clustered by subsidiary; the table reports β .

Table A13 asks the same question of the price increase an enrollee could not have anticipated from the block’s own rate schedule. Both brand groups’ closed blocks reprice well beyond never-closing blocks, and the dissimilar group’s point estimate is larger, but the two groups’ confidence intervals overlap, so the apparent contrast is not a finding.

What the null does and does not settle. There is no firm-side evidence that parents select dissimilar-named subsidiaries to execute deadpooling. Two readings survive. The mechanism may not operate through naming at all—a parent needs only that the enrollee fail to shop, not that they fail to recognise a name. Or the measure may be too coarse to capture the kind of distance that matters to a senior reading a rate letter. We report the null rather than resolve between them. Note also that an earlier version of this test, run on data through 2020, found a marginal effect in the predicted direction; it does not survive the extension to 2025, which is the reason it is reported here rather than in the text.

Table A12. Brand Distance Between Parent and Subsidiary Does Not Predict Deadpooling or Post-Closure Repricing

	All plans	Post-acquisition only
<i>Panel A: plan-level linear probability model</i>		
Deadpools	0.076 (0.061)	0.062 (0.081)
Ever closes	−0.050 (0.091)	−0.034 (0.112)
<i>Panel B: per-enrollee premium, dissimilar × post-closure</i>		
All plans	125.65* (73.41)	116.34 (82.81)
Closed plans	82.27 (72.13)	80.23 (87.89)
Deadpooling plans	39.76 (82.81)	48.14 (92.72)
Plans	17,364	13,796

Notes: All specifications absorb parent company, state and plan letter, are weighted by covered lives, and cluster standard errors at the subsidiary (NAIC) level; Panel B additionally absorbs calendar year. The regressor is an indicator for a subsidiary whose name shares no brand token with its parent's legal name, constructed as described in the text. Panel A collapses to one observation per plan, since both the outcome and the regressor are constant within a plan. Panel B is a plan-year panel and reports the coefficient on the interaction of the dissimilar-name indicator with an indicator for years at or after closure, controlling linearly for per-enrollee claims; the units are annual dollars per enrollee, inflation adjusted to January 2025. The right-hand column restricts to plans first observed at or after the year their subsidiary entered its current parent group, so that inherited pre-acquisition names cannot drive the result. Nothing in the table supports the brand-distance conjecture: within a parent, dissimilar-named subsidiaries are not detectably more likely to deadpool, are if anything less likely to close at all, and do not reprice their trapped pools more aggressively. The estimates are underpowered rather than precisely zero—the premium standard errors are of the same order as the coefficients—so a modest true effect would not be detectable here. * $p < 0.10$.

Table A13. Brand Distance Does Not Predict How Hard a Closed Block Is Repriced

	Unanticipated increase	Net of never-closing blocks	Deadpooled blocks
Brand-similar subsidiaries	\$583 (\$236)	\$459* (\$259)	270
Brand-dissimilar subsidiaries	\$857 (\$111)	\$733*** (\$166)	361
<i>Memo:</i> acquired subsidiaries	\$920 (\$47)	\$796*** (\$136)	231
<i>Reference:</i> never-closing blocks	\$124 (\$128)		2,130

Notes: The outcome is the five-year unanticipated price increase of Section 7.1: for an enrollee aged a in the year before closure (b), the CSG quoted price at age $a + 5$ five years later less the price the block's own schedule already showed for age $a + 5$ in year b , $p(a + 5, b + 5) - p(a + 5, b)$, so the anticipated age gradient is excluded. Entry ages 65–80 are pooled; annual January 2025 dollars, weighted by the block's enrollment in the baseline year. Never-closing blocks are measured over the same calendar window. The first column is the deadpooled blocks' mean; the second is that mean less the never-closing mean, estimated by regressing the increase on a deadpooling indicator among that group's deadpooled blocks and all never-closing blocks. Standard errors in parentheses are clustered by NAIC company, since each block enters once per entry age. Brand similarity is the subsidiary-level classification used above; six deadpooled blocks whose parent is the company itself have no brand distance and enter neither brand row. The acquired row is a memo item rather than a third category: it is a subsidiary whose parent identity changes during the panel, which overlaps both brand groups and can reflect a reporting reassignment as well as a sale. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

D.4 Deadpooling and Own Medicare Advantage Presence

Medicare Advantage is the main alternative to a Medigap policy, and most of the parents in our data sell both. That invites a natural question, which referees and seminar audiences have both raised: do insurers deadpool more aggressively in markets where they themselves have a larger MA presence, either because MA gives them somewhere to route the enrollees they price out, or because Medigap rents help fund MA expansion? We tested it and found nothing we are willing to report as a result. This appendix records the null so the question is not left open.

We build MA market share from the CMS State/County/Contract enrollment files, map contract-level organizations to their NAIC parent with a year-specific crosswalk, and collapse to the $\text{naicgroup} \times \text{state} \times \text{year}$ level, matching the unit at which we measure deadpooling. The outcome is the share of a parent's Medigap lives in that state-year held in deadpooling plans; the regressor is the same parent's MA market share in the same state-year. Parents with no MA enrollment in a state-year are retained at zero share. Standard errors are robust, and we run every specification both unweighted and weighted by the parent's Medigap lives, since the two answer different questions—the average *cell* against the average *enrollee*.

The estimate is not stable across those two weightings, which is the substance of the null. Pooling all parents, a percentage point more own-MA share is associated with 0.621 more percentage points of deadpooled lives unweighted (s.e. 0.242, $p = 0.010$, $N = 34,874$), but -0.243 once cells are weighted by the enrollees actually exposed (s.e. 0.566, $p = 0.667$). The sign reverses and the significance disappears. Dropping the paper's never-deadpooling benchmark carrier (Appendix C.11.1), which is also a large MA insurer, does not settle it: the unweighted estimate is 0.605 ($p = 0.012$) and the weighted one -0.929 ($p = 0.184$). We read it as an association carried by small cells rather than by the enrollees the mechanism operates on.

The deeper problem is that there is very little variation to identify anything from. Once that carrier is set aside, 99.6% of the remaining 34,058 parent $\times$ state $\times$ year cells have *zero* own-MA enrollment: the parents that dominate Medigap deadpooling are, with that one exception, largely not the parents selling MA in the same markets. Restricting to cells with any own-MA presence at all leaves $N = 141$, where the estimate is 0.358 unweighted ($p = 0.194$) and 0.887 weighted ($p = 0.057$)—neither significant at conventional levels, and on a sample far too thin to carry a claim.

Two limitations bound what this test could have shown even in principle. The MA data do not cover every parent in the Medigap panel (one substantial Medigap writer does not appear in them at all, and another exited MA partway through the sample) so the overlap is incomplete in a way that is not random. And the design is correlational: there is no source of exogenous variation in own-MA presence here, so even a stable coefficient would not have identified a cross-subsidy. We therefore treat the cross-market channel as untested rather than refuted, and nothing elsewhere in the paper rests on it.

D.5 Does Closure Look Different Where Enrollees Can Switch?

The closure event study of Section 6.1 pools all states. If deadpooling relies on enrollees being unable to leave, closure should look different where they can, so this appendix runs the same specification separately by regulatory regime, on the same deadpooling blocks as Table 2.

Table A14 reports the five-year ATT on each outcome for each regime. The contrast that identifies the question is between the forty-three states with no switching protections and the eight with an annual switching window. Where enrollees cannot leave, deadpooling closure raises the per-person premium by \$1,072 five years out and per-person claims by \$618, both significant, and the loss ratio changes by -0.073. Where an annual window exists, the per-person premium also rises significantly, by \$634, and the loss ratio changes by -0.074, so per-person profit is similar in the two regimes (\$454 and \$522). Two outcomes differ. Claims rise by \$112 where a window exists, which is not distinguishable from zero. And the quoted price, which pool composition cannot move, rises by \$779 where enrollees cannot leave but by \$373, not distinguishable from zero, where a window exists. Figure 8 shows that the practice is more common where switching is restricted; conditional on a deadpooling closure, the realized profit per enrollee is similar across these two regimes, but the repricing of the plan itself is concentrated where enrollees cannot switch.

Enrollment falls further where a window exists. Measured in logs, so that it is comparable across regimes with very different block sizes, enrollment in a deadpooling block changes by -60.3% within five years where switching is restricted and by -78.2% where a switching window exists. Together with the claims result, this suggests that a switching window changes who is left in a closed block: more enrollees leave, and the remaining pool does not become measurably sicker. These per-regime estimates are less precise than the pooled ones, and some pre-trend tests reject: premiums ($p = 0.001$), claims ($p = 0.016$) and enrollment ($p = 0.003$) where switching is restricted, and profit ($p = 0.025$) where a window exists.


The third regime, the three states operating a continuously open market, is reported for completeness rather than as a third comparison. Connecticut, Maine and New York are the states whose pre-existing rules make them unusable as controls in Section 8.2, and with three treated clusters the confidence intervals there should be read as indicative. Its profit estimate is positive but imprecise, its profit and loss-ratio pre-trend tests reject ($p = 0.000$ and $p = 0.000$), and we would not rest anything on it.

The per-person premium in Table A14 is a realized per-enrollee amount, so it carries pool composition as well as price. Figure A8 gives the quoted-price counterpart regime by regime: the same closure event study run on a fixed basket of ages, which composition cannot move. The ordering across regimes is the one the lock-in mechanism predicts, and reading each regime against its per-person premium in Table A14 separates what the carrier charged from who was left to charge it to.

Table A14. Deadpooling Closure Effects by Switching Regime


Five-year ATT	No switching protections (43 states)	Annual switching window (8 states)	Continuously open market (3 states)
CSG quoted-price index (\$)	779 (155)	373 (200)	-335 [†] (31)
Per-person premium (\$)	1,072 (239)	634 (252)	258 (204)
Per-person claims (\$)	618 (146)	112 (142)	-113 (116)
Per-person profit (\$)	454 (99)	522 (171)	371 (270)
Loss ratio	-0.073 (0.020)	-0.074 (0.027)	0.004 (0.081)
Enrollment (% change)	-60.3 (18.6)	-78.2 (30.2)	-78.3 (11.0)

Notes: Standard errors in parentheses, backed out of the 95% confidence intervals. [†]No continuously-open block is observed in the CSG quotes five years after closure, so this cell is the effect in year 3 after closure, the latest year estimated; its pre-closure coefficients are far from zero, so it should not be read as a causal estimate. The first row is the ATT on the CSG quoted-price index—a fixed (Laspeyres) basket over ages 65–85 weighted by the MCBS age distribution of Medigap holders, extended past the end of a closed block’s quotes by the closed-block rate increases the carrier filed—estimated on the same specification with control block-years restricted to blocks quoted in their own right that year. Because a fixed basket cannot move with the pool, it separates repricing from composition within each regime; the row below it is the realized per-enrollee premium. The by-regime cells are thin and the continuously-open column rests on three states: on the quoted-price index only 170 treated block-years survive there and no five-year cell is estimated at all, which is why that entry is an em-dash rather than a number. Read the other two columns as the comparison. Callaway–Sant’Anna ATT at event time +5, by state regulatory regime, from the specification of Figure 3 estimated separately within each regime. As in Table 2, the treated blocks are deadpooling blocks, and never-closing, never-deadpooling blocks are the control group in each. Dollar outcomes are per enrollee in 2025 dollars and are enrollment-weighted. Enrollment is the ATT on *log* enrollees, reported as a percentage change: mean pre-closure block size differs across regimes, so a level effect is not comparable across them. The three-state column rests on three treated clusters and its intervals are indicative rather than well-sized; those states are excluded as controls elsewhere in this section for the same pre-existing-rules reason.

Figure A8. The Quoted Price After Closure, by Switching Regime

Notes: Callaway–Sant’Anna ATT of deadpooling closure on the CSG quoted-price index (a fixed basket of ages 65–85 weighted by the MCBS age distribution of Medigap holders, annual 2025 dollars), estimated separately for each switching regime, against never-closing blocks in the same regime; doubly robust (*drip*) with state and plan-letter controls, universal base period (*long2*), enrollment-weighted; bars are 95% confidence intervals from NAIC-clustered standard errors. No continuously-open block is observed past three years after closure.

Figure A9. Premiums, Claims, and Profits Are Nearly Identical Across Switching Regimes



Notes: This plot is weighted by total lives and is calculated on block-years through 2020, matching the five-year look-forward window used to classify deadpooling elsewhere in the paper and the sample of Figure 8 beside it. Differences between switching-regime and unprotected states are essentially zero on all three measures, which is inconsistent with the claim that switching rules generate substantial adverse selection. Premiums and claims are higher in open-enrollment states, which is consistent with selection into the market under those rules—though as averages these cannot rule out offsetting reallocations within the market.

E Appendix: Consumer Harm: Alternative Measures and Responses

E.1 Alternative Measures of the Transfer

Section 6.1 measures the transfer as the premium rise beyond what the higher claims cost at the pre-closure loss ratio. Table A15 compares three routes to this number. The loss-ratio event study implies a markup widening of $(C_0 + \Delta C)(1/(LR_0 + \Delta LR) - 1/LR_0) = \$236$ per enrollee-year, or \$1,673 million. The CSG quoted prices, measured against always-open plans of the same letter in the same state and year, give \$376 per enrollee-year, or \$1,389 million for the 2014–2020 closures CSG quotes. We lead with the NAIC split for two reasons. First, the transfer is the insurer’s gain, premiums net of claims, and only the NAIC filings report both for the same blocks and years; a CSG quote is a price with no claims. Second, the NAIC filings cover every closure, while CSG quotes start in 2014, so the CSG route covers only the 1,182 blocks that close from 2014 on with their own quote, and none of the large 2010 cohort. The CSG route is the check rather than the headline, but it has one strength the NAIC routes lack: it compares the same age in both prices, so an aging pool cannot move it.

Table A15. Three Routes to the Transfer

Route	What it measures	Per enrollee-year	Total (\$M)
NAIC split (headline)	premium rise beyond claims at LR_0	\$243	1,723
NAIC loss-ratio drop	the same, from the loss-ratio ATT	\$236	1,673
CSG quotes	quoted price above always-open plans	\$376	1,389

Notes: All deadpooling closures 2006–2020, each followed for five years, 7,099,058 enrollee-years, January 2025 dollars. The NAIC routes use the column 3 premium, claims and loss-ratio ATTs of Table 2 (equation 8); per-enrollee NAIC figures divide by December-31 enrollment. The CSG route is the excess quoted premium of stayers over always-open plans of the same letter in the same state and year, for the 2014–2020 closures with their own quote, so its total covers fewer closures than the NAIC routes.

Valuing leavers on the realized premium. The main text values each enrollee who leaves Medigap at the Callaway–Sant’Anna estimate of the closure effect on the quoted price, the unexpected change in the price they face. An alternative values them on the realized NAIC premium path: the growth in the average premium a closed deadpooling block collects per enrollee beyond the growth of open blocks. That path also rises as the remaining enrollees age along their rate schedules and as the healthier leave, neither of which changes the price any one leaver was charged. On this path the loss from leaving Medigap is \$36.1 million per year of closures (\$541 million over all closures), against \$43.0 million on the Callaway–Sant’Anna path, and total lost consumer surplus is \$101.4 million per year of closures, or \$1.57 per dollar of insurer markup, against \$108.3 million and \$1.67. Over seven years the realized path gives \$163.8 million per year of closures and \$1.60 per dollar.

E.2 The Entry Discount That Would Compensate a Healthy Entrant

A price at 65 is not a lifetime price: the healthy entrant who buys the successor is the future trapped stayer. We therefore price the whole path. A healthy 65-year-old buys into an open block drawn from the enrollment-weighted distribution of block ages. Each year the block closes with a hazard estimated from the NAIC panel by block age and weighted by enrollment, so it is the hazard the average *enrollee* faces rather than the average block. On closure the enrollee escapes to another open block at the market price if they can still pass underwriting, and is otherwise trapped and rides the ramp. Escape probabilities are the age-specific CSG-rule pass rates of Appendix B (49.6% at 65–74, 33.1% at 75+), mortality is the age-band schedule of the MCBS pool, and premiums are discounted at 0.03.

Two modeling choices carry the result. First, escape is drawn *once*, at the moment of closure, not re-drawn each year: failing underwriting is a persistent personal characteristic, not an annual lottery, and treating it as iid lets almost everyone escape within a few years and drives the estimated harm to nearly zero (\$64 rather than \$1,710). Second, the ramp is the aging-netted wedge of Appendix C.10 rather than the raw one, since a closed block’s pool ages while an open block refreshes and part of the measured ramp is therefore the age curve rather than extraction.

A healthy entrant expects to pay \$1,710 more in present value over their remaining lifetime because deadpooling exists, against \$38,807 in the no-deadpooling counterfactual. Because every deadpooled premium scales with the entry price, an entry discount d scales the whole path, and there is a discount at which the entrant is exactly indifferent:


$$d^* = \frac{-PV_{\text{excess}}}{PV_{\text{cf}} + PV_{\text{excess}}}.$$

This is scale-free (the entry premium cancels) so it does not depend on how that premium is measured. Centrally $d^* = -4.2\%$: a successor would have to enter four percent below market to compensate a healthy 65-year-old for the closure risk they bear. The estimated entry discount comes from CSG quotes in each new block’s opening year. We regress the log quoted rate for a 65-year-old on indicators for a successor block (a fresh block opened by a deadpooling parent within five years of one of its closures)

and for a block that later closes, with state $\times$ plan $\times$ year fixed effects so that the comparison is with rivals selling the same plan in the same state and year, and standard errors clustered by parent group. A successor enters 2.5% above the market (se 2.3, $N = 2,420$), with a confidence interval of $[-2.0, +7.0]$, entirely above that threshold. Across the grid of wedges and escape probabilities, every value of d^* is negative, so the healthy require some discount everywhere in the parameter space; the most generous cell available (the aging-immune wedge combined with a 95% escape probability, the rate the model itself assigns to its healthy consumers) still gives $d^* = -0.8\%$ and \$299 of expected excess. Under no parameterization does the model's assumed 48.5% discount (in the model of Section 9, successor blocks open at \$1,488 against a market-average premium of \$2,862) approach break-even.

The successor's entry price is not low in the age cross-section either. Figure A10 plots the quoted premium by attained age: deadpooling plans quote more than other plans at every age, including 65.

Figure A10. The Same Age Profile in Dollars: Deadpooling Plans Quote More at Every Attained Age




Notes: Mean quoted premium by attained age, deadpooling plans against other plans: CSG quoted premiums for a female, non-tobacco applicant, attained-age plans, 2014–2025, residualized on state $\times$ plan $\times$ year $\times$ age fixed effects. The age-specific mean log rate is added back onto the residual and the sum exponentiated, which restores the level of the age profile while keeping the market $\times$ age netting that makes the two series comparable. Each point is therefore a geometric mean annual rate at that attained age, net of which markets each group sells in, annualized from the monthly quote and deflated to January 2025 dollars. Bands are the exponentiated log-scale 95% interval and are narrow enough to be hard to see at this scale.

The same holds for the average offer, not only the cheapest ones. Figure A11 repeats Figure 6 with the mean quoted rate at 65 in place of the tenth percentile. Moving a market from no deadpooling to all of it raises the mean quote by 9.7 log points (se 4.2, $p = 0.025$), against 10.6 for the tenth percentile.

E.3 Enrollees Leave in Response to Premium Increases

A central premise of the deadpooling mechanism is that policyholders leave a closed block *because* they are charged rising premiums, not simply dying off or exiting at random. This appendix examines that premise directly. For each deadpooling block we measure first-year disenrollment as $1 - L_{t_c+1}/L_{t_c}$,

Figure A11. The Average Quote Is Also No Lower Where Deadpooling Is Prevalent




Notes: As Figure 6, with the mean quoted rate at 65 (in logs) as the outcome. Market = state $\times$ plan letter $\times$ year. Both axes are residualized on market fixed effects, plan $\times$ year fixed effects, and the number of offers on the shelf. Binned means, twenty equal-size bins, over markets in all states. Descriptive: deadpooling intensity is not randomly assigned. The coefficient in the text is from the regression on all markets with the offer-count control, standard errors clustered by state.

where t_c is the year the block closes to new enrollment and L is enrollment, so the window runs from the closed pool forward one year. We relate it to three measures of the price policyholders face over the same window: the markup *level* (premium/claims), premium *growth* ($p_{t_c+1}/p_{t_c} - 1$), and the absolute premium *increase* in dollars. All statistics are enrollment-weighted and variables are winsorized at the 1st and 99th percentiles; fitted lines are OLS. The sample is the 2,356 deadpooling closures with observed enrollment and premiums in both the closure year and the following year.

Figure A12 gives the aggregate version of the price relationship, binning closing plans into deciles of premium change and plotting the associated enrollment loss; it is the figure the main text points to when it notes that beneficiaries respond to price. The block-level regressions that follow decompose that binned pattern by the specific price measure enrollees face.

Disenrollment tracks how much the premium rises. The first-year relationships in this appendix use the realized premium per enrollee, which has enrollment in its denominator: when enrollees leave, the premium per remaining enrollee can move for that reason alone, so a slope estimated on it may be partly mechanical. The quoted price has no such link to enrollment. Across all blocks with a quoted price, each additional \$100 of quoted price changes a block's enrollment the following year by -5.20 percent (s.e. 0.76), against -1.90 percent (s.e. 0.18) per \$100 of realized premium on the same blocks, both with block and year fixed effects (Figure A13). The response to the quoted price is the larger of the two, so the mechanical link does not produce it, and it is the estimate Section 7.3 uses. Within closing deadpooling blocks, the realized-premium version points the same way: each additional \$100 of first-year premium increase is associated with 1.2 percentage points more disenrollment (Figure A14); in percentage terms, a


Figure A12. Larger Premium Increases Are Associated with Larger Enrollment Losses



Notes: This plot scatters deciles of the one-year change in per-enrollee premium over the closure year (January 2025 dollars) against the percent decrease in enrollees each plan experiences over the same year. The window is one year to limit the role of mortality. Deciles are defined using total lives as weights. Closures are defined based on calendar year, so a closed plan gaining lives (as displayed in the rightmost point) would have made those gains in the part of the year before closure. Closure dates cluster heavily on a semiannual convention: among the 7,717 plans with a recorded closure date, 30% close effective the end of May or June 1 and 26% the end of December or January 1, so 56% fall on one of these two dates. The upward slope is consistent with enrollees responding to price when deciding whether to disenroll: plans in the top decile of premium increases (about \$1,000 a year) lose roughly 30% of their enrollment. The relationship is descriptive rather than causal (premium growth and disenrollment are jointly determined, since the departure of healthier members itself worsens the pool and pushes premiums up) and which margin leavers exit to (another Medigap, Medicare Advantage, or no supplemental coverage) is not observable in these data.


block whose premium rises ten percentage points faster loses roughly an additional 2.7 percentage points of its enrollment in the first year (pooled slope 0.27, s.e. 0.07; 0.36 with state and block-age fixed effects; not shown). The markup *level* (premium over claims), which enrollees never observe, has no relationship with exit (slope 0.001, s.e. 0.008; not shown).

Figure A13. Enrollment Falls When the Quoted Price Rises



Notes: Binned means (twenty equal-size bins) of the one-year percent change in a block's enrollment against the one-year change in each premium measure, per \$100, with OLS fits. Squares: the CSG quoted-price index, a price schedule for a fixed basket of ages, computed without reference to who is enrolled. Circles: realized premium per enrollee from the NAIC filings, whose denominator is enrollment. All blocks with a quoted-price index; the extreme 1% of each variable is trimmed. Not enrollment-weighted, because the outcome is enrollment. With block and year fixed effects and standard errors clustered by NAIC company, the slopes are -5.20 (s.e. 0.76; $N = 41,745$) for the quoted price and -1.90 (s.e. 0.18; $N = 41,974$) for the realized premium.


Figure A14. In Closing Blocks, Disenrollment Rises With the First-Year Premium Increase



Notes: First-year disenrollment against the first-year per-person premium increase, with the horizontal axis labelled in dollars. Enrollment-weighted, winsorized at the 1st/99th percentiles. The pooled slope is 0.012 per \$100 (s.e. 0.003). Premium increases are the common case rather than the exception over this window: 87.9% of enrollees in a closing block face a premium that rises between the closure year and the following year. The dollar-denominated version reproduces the percentage result reported in the text, so the response is not an artifact of scaling by the base premium. The premium here is realized premium per enrollee, whose denominator is enrollment; Figure A13 gives the version on the quoted price, which has no such link.

Ruling out mortality. The response is not an artifact of sicker, older pools dying off. First, the premium-growth relationship holds within strata of two separate proxies for a block’s mortality risk. Splitting on block age (years since the block’s filed approval, a proxy for the pool’s age) the slope is positive in all three terciles (0.31, 0.35, and 0.53) and significant in the youngest, lowest-mortality blocks ($p = 0.002$) as well as the oldest; the middle tercile is imprecisely estimated (standard error 0.22) rather than flat (Figure A15). Splitting instead on pre-closure per-person claims (a proxy for pool health) the slope is 0.35 in the lowest-claims/low-mortality quartile and 0.44 in the highest-claims quartile (Figure A16). Second, a direct accounting bounds the mortality contribution: applying Social Security period life-table death rates (Social Security Administration 2024) to the MCBS age mix of the Medigap population implies expected annual mortality of about 3.9%, against mean first-year disenrollment of 17.4%, so roughly 78% of first-year exits are not attributable to death (and 43% even under a deliberately high 10% mortality ceiling). Random attrition is likewise inconsistent with the evidence, since idiosyncratic exit would not co-move with premium growth.

Figure A15. The Premium Response Holds Even in the Youngest, Lowest-Mortality Blocks




Notes: Binned scatter with OLS fits by tercile of block age, a proxy for pool age and hence mortality risk. Block age is dated from the block’s filed approval date, not from its first year in the panel: the panel begins in 2006, so dating from it would floor every earlier-approved block at 2006 and understate the age of the 40.4% of the 2,203 blocks here whose approval predates the panel, pushing genuinely old pools into the youngest tercile. Enrollment-weighted. The slope is positive in all three terciles (0.31, 0.35, 0.53) and significant in the youngest ($p = 0.002$) and the oldest; the middle tercile is imprecisely estimated (standard error 0.22) rather than flat. The response holds among the youngest, lowest-mortality blocks, so the disenrollment response cannot be death in disguise.

Cost pass-through versus a price response. A remaining concern is that premiums might be rising only to track the block’s own worsening costs, so that “response to premiums” is really a response to rising claims. Controlling for contemporaneous claims growth, however, the premium-growth coefficient survives and if anything strengthens (to 0.32, $p < 0.001$), while claims growth on its own carries no independent relationship with exit (-0.06 , not significant). Figure A17 plots the partialled-out (Frisch–Waugh) relationship: disenrollment and premium growth after each is purged of its claims-growth component. Policyholders leave when the premium rises, not merely when the underlying cost of their pool rises.


The two coefficients are worth reading together, because each rules out a different alternative and only

Figure A16. The Premium Response Holds in the Healthiest Blocks



Notes: Binned scatter with OLS fits for the extreme quartiles of pre-closure per-person claims, a second proxy for pool health and mortality risk. Enrollment-weighted. The response is present among the healthiest blocks as well as the sickest: the slope is 0.35 (s.e. 0.10) in the lowest-claims quartile against 0.44 (s.e. 0.06) in the highest. The two interior quartiles are not plotted because they are uninformative rather than genuinely flat: the second quartile's slope is 0.12 (s.e. 0.10) and the third's is 0.12 (s.e. 0.48), the latter estimated so imprecisely that its confidence interval spans both extremes. That the response survives where mortality risk is lowest is what a mortality explanation cannot produce.

Figure A17. Enrollees Leave When the Premium Rises, Not When Costs Rise



Notes: Frisch–Waugh residualized scatter: first-year disenrollment and first-year premium growth are each regressed on first-year claims growth, and the residuals plotted against one another. The slope (0.32) equals the premium-growth coefficient in the multivariate regression that also includes claims growth. Enrollment-weighted; marker size proportional to enrollment. Controlling for the block's own claims growth, the premium-growth coefficient survives and strengthens to 0.32 ($p < 0.001$), while claims growth carries no independent relationship with exit (-0.06 , insignificant). The response is to the price, not to the underlying cost of the pool.

the pair settles the question. That premium growth predicts exit *holding claims fixed* rules out the health story: if enrollees were leaving because they were getting sicker or dying, the claims term would carry that and the premium term would not survive its inclusion. It does survive, and strengthens. That claims growth carries no independent relationship *holding premiums fixed* rules out the converse: rising costs do not move people except through the price they are charged. So price is the operative margin, and the exit we observe is a behavioral response to being repriced rather than a mechanical consequence of a deteriorating pool. This is a statement about the *margin* people respond to, not a claim that selection is absent—who leaves is still selected on health, which is the composition effect the closed-block claims decomposition addresses separately.


Taken together, these patterns show that disenrollment from deadpooling blocks tracks the premium *increase* (in percentage or dollar terms) rather than the premium level, that this holds within low-mortality subpopulations and after netting out expected deaths, and that it reflects the premium itself rather than cost pass-through. We interpret this as descriptive evidence consistent with a demand response and inconsistent with the pure-mortality and pure-attrition alternatives. It is not a causal estimate: premium growth and voluntary exit are jointly determined, since the departure of healthier members itself worsens the pool and pushes premiums up. The responsiveness of exit to price is instead a calibrated parameter of the model in Section 9.

E.4 Responsiveness to Premium Increases Under Switching Rules

Equation (14) compares closed deadpooled blocks in the same state, plan type, and year that faced increases of different sizes. Adding block fixed effects α_i instead compares each block only with itself in years with smaller or larger increases. Figure A18 reports this version. Without a switching rule, the response is now close to flat: relative to a year with no increase or a decrease, a block loses an additional -0.2 , 0.5 , and 1.0 percentage points of its enrollees after increases of 0–5%, 5–10%, and 10% or more. With a switching rule in force, it loses 7.4, 10.3, and 7.5 percentage points, and the difference from the no-rule response is significant at the 1% level in all three bins.

A linear version of equation (14), which replaces the bins with the continuous increase x_{it} and the outcome with the change in log enrollment, gives the same picture and allows randomization inference over the adoption years (200 draws). In this version a 10% unexpected increase is followed by the loss of 0.8% of enrollees without a switching rule and 6.0% with one. The interaction (-0.56 , randomization-inference $p = 0.005$) does not rest on comparing the early adopters with other states: identified only from the change within the four 2022 adopters it is -0.60 ($p = 0.025$) enrollment-weighted and -0.38 ($p = 0.209$) unweighted, although the unweighted estimate pooling all adopters is indistinguishable from zero (0.15 , $p = 0.512$). Closed blocks that did not deadpool show no such response (0.21 , $p = 0.333$), though the difference between the two kinds of closed block is itself imprecise (-0.21 , $p = 0.234$). Carrier A, the large deadpooling carrier of Section 8.2, accounts for 59% of the enrollment-weighted treated sample, and dropping it strengthens the interaction to -1.18 ($p = 0.005$).

Figure A18. Responsiveness to Premium Increases with Block Fixed Effects




Notes: As Figure 11, adding block fixed effects α_i to equation (14). Grey bars plot β_k ; blue bars plot $\beta_k + \gamma_k$, each relative to a year with no increase or a decrease. Observations are weighted by the prior year's enrollment; 95% intervals use state-clustered standard errors. $N = 6,520$ block-years; blocks observed in only one year drop out once block fixed effects are included.

E.5 Enrollees Place Little Value on Their Insurer

If enrollees valued staying with a particular insurer, the brand could discipline the repricing of a closed block. They do not. Our Understanding America Study module asks Medicare enrollees whether it would matter which company insures them and, if so, how much more they would pay each year to have a plan issued by their preferred company rather than identical coverage from the cheapest seller. Among 1,849 respondents (survey-weighted), 83% have no preferred insurer or would pay nothing extra to stay with it (66% have no preference; 16% have one but would pay nothing). 96% would pay less a year than the \$435 premium rise a deadpooled block imposes, and 96% less than the \$398 profit rise. Figure A19 shows the distribution. Almost no one would pay a brand premium anywhere near the repricing.

Figure A19. Willingness to Pay to Stay with a Preferred Insurer



Notes: Understanding America Study, Medicare enrollees, survey weights; $N = 1,849$. Respondents with no preferred insurer count as zero; 23 with a preference but no answer are dropped. The right bar compares each answer with the deadpooling closure premium ATT (Table 2, column 3).

E.6 What Lock-In Buys: Stable Pools and Who Enters Them

Everything above prices what lock-in costs. It is worth asking what it buys. The defense of underwriting is that it keeps risk pools intact: if the healthy cannot be skimmed away and the sick cannot pile in, a block should be able to hold its price steady and stay open. This subsection tests that claim on its own terms, and finds that it is partly true—which is the strongest case the practice has.

How many pools are actually stable? Call a pool *stable* if it never closes, stays open at least five years, and never raises its price by more than a given amount from one year to the next. Measured on *realized* per-enrollee premiums, almost none is: only 1.5% of enrollees sit in a pool that clears a ten percent bar. That figure is misleading, and the reason matters. Realized per-enrollee premium is total premium over covered lives, so it moves with who is enrolled as much as with what is charged—a block whose healthy members leave posts a “spike” none of its remaining enrollees was billed. The quoted rates price the same hypothetical applicant every year, so composition cannot move them, and they are the right instrument here for exactly the reason they were in Section 7.2.

On quoted rates the picture reverses. Of 12,301 blocks in the CSG panel, 4,083 never close and stay open at least five years, and among all life-years 15.7% sit in a pool whose rate for a 65-year-old never rises more than ten percent in a year; at 75 the figure is 16.8%. Tighten the bar to five percent (below ordinary medical trend) and stability becomes rare again, 2.0% and 2.2%. So most pools do not reprice violently. The instability the realized series shows is overwhelmingly *composition*, not repricing, and that is itself a finding: the per-enrollee premium series that the literature typically observes overstates how much price risk an enrollee bears.

Stability is highest where switching is most restricted. Splitting by regime, 18.5% of life-years in states with no switching protections sit in a stable pool at age 65, against 5.4% under an annual switching window and 2.5% under open enrollment; at 75 the gap is wider still (19.8% against 2.5%). The ordering runs the way the stability defense predicts, and it is the one cut in this paper where restricted switching looks good.

Who enters a new block, and where. The mechanism behind that ordering is visible at entry. A block opening in a state where underwriting binds can screen its applicants; one opening where anyone may buy in annually cannot. We score every block that opens after the panel begins on its first-year per-enrollee claims, relative to the lives-weighted claims of every block in the same state, plan letter and year—a within-cell benchmark, so a state’s cost level cannot drive the comparison. Across 12,255 new blocks, the median entrant opens at 0.22 of incumbent claims where there are no protections and 0.26 under an annual switching window, but 0.43 under open enrollment; weighting by the entrant’s own enrollment the contrast is 0.51 and 0.39 against 0.90. Where buyers can enter without underwriting, a fresh block opens with a pool that looks like the market; where they cannot, it opens at half the market’s risk or better.

Reading this against the rest of the paper. The same screening that makes deadpooling profitable—the underwriting that keeps a closed block’s sick from leaving —also keeps a new block’s pool from filling with them. Lock-in is not doing nothing. It buys measurably steadier prices and measurably cleaner entry, and a market design that removes it should expect to give some of that up. What this paper documents is the price of that purchase: the same mechanism hands the incumbent insurer a captive pool to reprice, and Sections 7.1 through 7.3 measure what it extracts from it. The defense of lock-in is real but partial, and the question a regulator faces is not whether it works but whether it is worth what it costs.

Three caveats travel with these numbers. The open-enrollment cell rests on 188 entering blocks in three states, so it is indicative rather than precise. The quoted-rate panel runs 2014–2025, a shorter window than the claims panel, so the five-year screen is judged over twelve years rather than twenty. And the entry ratio measures a block’s *ability to screen*, not welfare: a new block that cannot skim is evidence that a protection binds, not that the market is worse off for it.

F Appendix: The Medicare Agent Survey

The NAIC and CSG data describe deadpooling from the insurer’s side. To see how the intermediaries who actually place seniors into plans perceive and respond to it, we fielded a survey of licensed Medicare insurance agents. This appendix describes that instrument and reports its descriptive results; the survey is not used to identify any of the paper’s causal estimates and is offered as corroborating supply-side evidence.

Data collected. The Medicare Agent Survey is a pilot fielded in 2026 through the Cint online research panel, screening for respondents who hold an active Medicare insurance license. We describe only the Medigap-relevant content here. A deadpooling module measures awareness of block closures, the perceived share of clients trapped on a closed block, beliefs about the underwriting outcomes those clients face, how often an agent weighs a carrier’s closing history, per-carrier suspicion that a company actively deadpools, and what an agent would advise a client facing a large rate increase. A three-item knowledge screen tests whether agents know that switching from Medicare Advantage to Medigap later requires medical underwriting, that Original Medicare alone has no annual out-of-pocket maximum, and that same-letter Medigap benefits are standardized by law. The instrument also carries modules on agent demographics and book composition, client vignettes eliciting Medicare Advantage versus Medigap recommendations, a commission-gradient experiment, and existing-client service, which we describe no further because the paper does not draw on them.

Automated-response screening. Online agent panels are known to attract automated and low-effort responses, so the instrument embeds several traps and we define a graduated set of exclusion screens. Two honeypots catch page-scraping bots: an instruction rendered in invisible text that asks the respondent to include a specific phrase (a human never sees it; a bot that reads the page source often

complies), and an off-screen link that only an automated agent would click. A third trap hides an instruction to write an implausibly long open-ended answer. Beyond these deterministic traps we flag near-verbatim open-ended answers shared across distinct respondents (a signature of answer farms and language-model output that within-respondent duplicate checks miss) as well as attention-check failures, the knowledge screen above, a reCAPTCHA bot score, and duplicate device and network fingerprints. Of 1,044 responses, 627 were complete. We define four nested analysis samples of completed responses: automated/bot responses removed only ($N = 614$); additionally excluding attention-check failures ($N = 476$); instead excluding failures of the three-item knowledge screen ($N = 410$); and a strict screen that removes attention, knowledge, straightlining, and low bot-score responses ($N = 366$). The figures below use agents who pass the attention and invisible-text checks, and each figure note gives the number of agents who answered its item.


What the survey shows. Two patterns recur. First, agents are broadly aware of deadpooling and its consequences: a majority recognize block closures, and they estimate that a large minority of trapped clients are either offered underwritten Medigap only at uncompetitive rates or denied it entirely (Figures A20–A22). Second, agents’ suspicion of specific firms is only loosely tied to which firms actually deadpool: the firms the NAIC data identify as deadpoolers draw suspicion at rates similar to firms that do not, including the market’s largest never-deadpooling operator (Figure A24). We therefore read the firm-level figures as measuring the market’s *perceptions* rather than as independent evidence on which firms deadpool, which is the task of the NAIC event studies. We also caution that the modules shown only to post-pilot respondents have small denominators, so the figures for those modules are directional rather than precise.

G Appendix: The Model

G.1 Market Structure


The market consists of $F = 37$ parent insurance companies, each operating across $S = 5$ states and offering $L = 3$ standardized plan tiers—three representative plans calibrated to aggregate market moments rather than models of Plans F, G, and N specifically. Each company–state–plan-tier triple defines a *block*: a pool of enrollees sharing a community-rated premium. Blocks are indexed by $b \in \mathcal{B}$. There are thus up to $F \times S \times L = 555$ potential blocks active at any time, though not all are open. Each block b is characterized by its current premium P_b , its open/closed status, its age, and its risk-pool composition. Community rating requires all enrollees in a block to pay the same annual premium, regardless of health status or age. New enrollees may only enter open blocks.

Figure A20. Most Agents Are Aware of Deadpooling / Block Closures




Notes: Share of agents aware that Medigap blocks close to new enrollment while continuing to insure incumbents. Agents who pass the survey's attention and invisible-text checks; $N = 305$ agents answered this item. Question asked: "Are you familiar with the practice of 'deadpooling', 'block exits', or 'subsidiary closures' where insurance companies close plans, concentrate high-cost enrollees together, raise premiums, and use subsidiary companies to open new books of business?"

Figure A21. Agents' Estimates of the Share of Clients Trapped on a Closed Block




Notes: Distribution of each agent's estimate of the percentage of their Medigap clients currently stuck on a closed block. Agents who pass the survey's attention and invisible-text checks; $N = 238$ agents answered this item. Question asked: "Out of every 100 of your clients with a Medigap policy, about how many are stuck on a closed block paying above-market premiums?"

Figure A22. Perceived Underwriting Outcomes for Clients When a Block Closes




Notes: Mean share of clients an agent places in each underwriting outcome (offered underwritten Medigap at a competitive rate, at an uncompetitive rate, not offered, or no belief), elicited as a constant-sum battery. Agents who pass the survey's attention and invisible-text checks; $N = 305$ agents answered this item. Question asked: "Think about all of your current Medicare clients. If all of your clients were to apply for a Medigap plan, about how many would fall into each of the categories below? (The four estimates should sum to 100%.)"

Figure A23. How Often Agents Weigh a Carrier's Block-Closing History




Notes: Share of agents choosing each frequency. Agents who pass the survey's attention and invisible-text checks; $N = 305$ agents answered this item. Question asked: "When recommending a Medigap plan, how often do you consider whether a carrier has a history of closing blocks and reopening new blocks under subsidiaries?"

Figure A24. Agents' Suspicion of Each Carrier Against Its Actual Deadpooling



Notes: Bars: share of agents rating each carrier as one that “actively deadpools” (rather than trusting it to keep blocks open, or unsure). Maroon: share of that carrier’s Medigap enrollees (life-years) in a plan the NAIC data classify as deadpooling, filings through 2020. Carriers are anonymized; * marks an active deadpooler in the NAIC data; the benchmark is the largest carrier that never deadpools. Agents who pass the survey’s attention and invisible-text checks; $N = 305$ agents answered this item. Question asked: “For each carrier below, indicate whether you believe the carrier actively engages in ‘deadpooling’ (closing Medigap blocks and reopening under new subsidiaries) or the carrier can be trusted to keep its blocks open.”

Figure A25. What Agents Would Advise a Medigap Client Facing a Large Rate Hike



Notes: Share of agents choosing each recommendation. Agents who pass the survey’s attention and invisible-text checks; $N = 305$ agents answered this item. Question asked: “Suppose the following client receives a 15%+ premium increase on their current Medigap plan. What would you recommend if the client is... a healthy 72-year-old client.”

G.2 Consumer Demographics and Health

G.2.1 Aging and Mortality

Each consumer i enters the model at age 65 and ages one year per period. Mortality depends on both age and health status:

$$\mu(a, h) = \mu_0^h \times m_\mu(a), \quad (19)$$

where μ_0^h is the age-65 baseline mortality rate for health state h and $m_\mu(a)$ is an age multiplier calibrated from MCBS and applied via linear interpolation between anchor ages. The baseline rates are $\mu_0^{\text{healthy}} = 0.015$, $\mu_0^{\text{moderate}} = 0.030$, and $\mu_0^{\text{sick}} = 0.080$. The moderate baseline is calibrated from self-assessed health relative-risk estimates in the MCBS: beneficiaries reporting “good” health have approximately twice the annual mortality of those reporting “excellent” or “very good” health (Lee et al. 2003), yielding $0.015 \times 2 = 0.030$. The mortality multipliers are $m_\mu(65) = 1.000$, $m_\mu(75) = 2.500$, and $m_\mu(85) = 6.667$, with linear extrapolation beyond age 85. These come from the *SSA 2021 period life table* (combined male and female, the arithmetic mean of q_x within each age bin), not from the MCBS: the MCBS public-use file records no death or vital status, so it cannot identify mortality at all. Age multipliers for expected claims are a separate matter and *are* from the MCBS, applied via linear interpolation between anchor ages, with values of 1.000 at age 65, 1.204 at age 75, and 1.299 at age 85. They are ratios of the mean count of reported conditions by age group rather than of dollar spending, since the public-use file carries health-burden outcomes and not claims; the condition-count ratio is a proxy for how costs rise with age, and a conservative one, as counts understate the convexity of spending in age. Beyond age 85, both multipliers are extrapolated linearly using the slope of the 75–85 interval.

G.2.2 Health States and Transitions

Each consumer occupies one of three health states: *healthy* ($h = H$), *moderate* ($h = M$), or *sick* ($h = S$). Annual medical claims are drawn from a normal distribution truncated at zero. The age-adjusted mean claims are $\bar{c}_H = \$1,050$, $\bar{c}_M = \$1,600$, and $\bar{c}_S = \$2,900$, with coefficients of variation $CV_H = 2.4$, $CV_M = 2.0$, and $CV_S = 1.8$ respectively. The relative ratios across states (1.00 : 1.54 : 2.77) are calibrated from the MCBS Cost Supplement PUF (2023), and the overall level is set so that the enrollment-weighted average claims matches the NAIC market average of \$2,255 per enrollee. New 65-year-olds are drawn healthy, moderate, or sick with probabilities 0.80, 0.15, and 0.05. That split is an assumption rather than an estimate: it is not calibrated to any moment and no MCBS analogue pins it, and it matters because it sets the starting composition of every fresh block and hence the level of the teaser price a successor can offer. The implied entry-cohort claims level, $0.80(1,050) + 0.15(1,600) + 0.05(2,900) \approx \$1,225$, is the figure the model uses for a newly opened pool.²⁸

28. The MCBS measure used is PAMTALPR (total private insurance payments), restricted to beneficiaries on traditional Medicare with no Medicaid and positive private insurance payments. CV estimates by chronic-condition group (0–1, 2–3, 4+ conditions) are 2.43, 2.00, and 1.81 respectively.

Health transitions follow a first-order Markov process. Health can only worsen:

$$\Pr(h_{t+1} = M \mid h_t = H, a) = p_{HM}^0 + \delta \cdot (a - 65), \quad (20)$$

$$\Pr(h_{t+1} = S \mid h_t = M, a) = p_{MS}^0 + \delta \cdot (a - 65), \quad (21)$$

Here $h_t \in \{H, M, S\}$ is the consumer's health state in year t (healthy, moderate, or sick) and a is their age. The baseline one-year transition probabilities are $p_{HM}^0 = 0.06$ from healthy to moderate and $p_{MS}^0 = 0.09$ from moderate to sick, and $\delta = 0.003$ is the amount by which each rises per year of age above 65. Consumers enter at age 65 and never earlier, so $a - 65 \geq 0$ always and no truncation at zero is needed. Once a consumer reaches the sick state they remain there permanently, so S is absorbing.

G.3 Consumer Demand

G.3.1 Willingness to Pay

A consumer's WTP for Medigap coverage is the certainty equivalent of the claims the coverage absorbs, computed by forward-simulating a 20-year Markov chain from the consumer's current health state. Let $\pi_\tau(h')$ denote the probability of being alive and in health state h' at horizon τ , and $S_\tau = \sum_{h'} \pi_\tau(h')$ the survival probability. The WTP center is the survival-weighted average annual claims plus a CARA risk premium proportional to the survival-weighted average annual claims variance:

$$\bar{v}_h = \frac{\sum_{\tau=1}^T \sum_{h'} \pi_\tau(h') \bar{c}_{h'} m_c(a_\tau)}{\sum_{\tau=1}^T S_\tau} + \frac{\rho}{2} \cdot \frac{\sum_{\tau=1}^T \sum_{h'} \pi_\tau(h') (\text{CV}_{h'} \bar{c}_{h'} m_c(a_\tau))^2}{\sum_{\tau=1}^T S_\tau}, \quad (22)$$

where $T = 20$ and ρ is the CARA absolute risk-aversion coefficient. Risk-averse consumers therefore value coverage above its expected cost; setting $\rho = 0$ recovers the actuarially-fair (expected-claims) WTP. With $\rho = 1 \times 10^{-4}$, this produces WTP centers of $\bar{v}_H = \$1,821$, $\bar{v}_M = \$2,716$, and $\bar{v}_S = \$3,580$ at age 65 (versus \$1,556, \$2,285, and \$3,199 at $\rho = 0$). Individual WTP is drawn from $\mathcal{N}(\bar{v}_{h_i}, (\sigma_v \bar{v}_{h_i})^2)$ truncated to $[0, 4\bar{v}_{h_i}]$ with $\sigma_v = 0.40$, and redrawn each year as health evolves.

G.3.2 Price Sensitivity

Each consumer i is assigned a price-sensitivity index $\lambda_i \in [0, 1]$ drawn from $\text{Beta}(\alpha \approx 0.60, \beta \approx 2.40)$, parameterized with mean $\mu_\lambda = 0.20$ and standard deviation $\sigma_\lambda = 0.20$. This is a right-skewed distribution with density decreasing monotonically from the origin, placing most mass at low price-sensitivity values, consistent with UAS survey evidence that 68% of enrollees consulted only one agent and 73% have never switched. Price sensitivity is drawn once at birth and held fixed throughout a consumer's life.

G.3.3 Enrollment Decision

Each year, $N_{\text{new}} = 7,000$ new 65-year-olds enter the market and enroll in Medigap. Entrants are allocated across open blocks with probability inversely proportional to premium. New entrants face no

underwriting—they have guaranteed issue at age 65. Underwriting applies only to mid-market switches and re-entry: $\rho_H = 0.95$, $\rho_M = 0.25$, $\rho_S = 0.05$, with an additional age penalty of 1 percentage point per year above 65.

G.3.4 Switching

Each period, enrolled consumers consider switching to a cheaper block. Consumer i shops with probability $\Pr(\text{shop}_i) = 0.20 \times \lambda_i$. Conditional on shopping, consumer i switches only if annual savings exceed a price-sensitivity-scaled threshold:

$$\text{switches} \iff P_{b_i} - P_b^{\min} \geq \frac{\bar{s}}{\lambda_i}, \quad (23)$$

where $\bar{s} = \$75$. The base shopping rate $p_s = 0.20$ and threshold $\bar{s} = \$75$ are calibrated jointly to reproduce the approximately 2% annual switching rate in the UAS survey. Switching also requires passing underwriting, which is where health enters the consumer's problem: a healthy applicant passes with probability 0.95, a moderate one with 0.25 and a sick one with 0.05, each falling a further point per year of age above 65, and the screen is waived inside a birthday-rule or guaranteed-issue window. Those probabilities are specified in the enrollment subsection above; a consumer who fails the screen stays where they are regardless of how large the saving would have been, which is the mechanism that makes the closed pool captive.

G.3.5 Exit and Re-entry

Consumers exit the Medigap market if premiums significantly exceed their WTP:

$$\Pr(\text{exit}_i) = \lambda_i \cdot \gamma^{\text{exit}} \cdot \max\left(0, \frac{P_{b_i} - v_i}{P_{b_i}}\right)^\eta, \quad (24)$$

where $\eta = 0.50$ governs the shape of the response (discussed below) and $\gamma^{\text{exit}} = 0.09$ is the exit responsiveness, calibrated to match the 0.56% annual rate at which Medigap enrollees leave the market, measured directly in the UAS module.

The shape of the exit response. Equation (24) raises the proportional excess of premium over willingness to pay to a power η , which the calibration searches jointly with the other parameters (Appendix G.9). With $\eta = 1$ the response is linear, so γ^{exit} scales every consumer's exit hazard by the same factor; $\eta > 1$ concentrates exit among the most heavily repriced, and $\eta < 1$ does the opposite, making the mildly repriced leave relatively more readily than the heavily repriced. A convex response is the natural candidate for the one moment the model does not reproduce: post-closure retention is too high at every horizon, which is a statement about *who* leaves rather than *how many*. The objective does not favor it. Over values of one and above, the linear rule fits best: convexity suppresses exit wherever the excess is small, which is most of the market, and the aggregate exit rate falls well below its target. That optimum sat on the edge of the grid, and extending the search below one moves it further, to the

concave $\eta = 0.50$ of the current calibration. The fit thus favors the aggregate exit rate over the shape of exit, consistent with the retention curve remaining too high.

We read this as evidence that the retention shortfall is not reachable through the functional form of the exit margin. The remaining candidate is the other route out of a closed block (passing underwriting to switch carrier) which is governed by the acceptance probabilities rather than by Equation (24).

The retention shortfall is an underwriting margin, not an exit margin. We checked this directly, varying the acceptance probability for enrollees in the intermediate health state, ρ_M , while holding every other calibrated parameter at its fitted value. Retention responds to ρ_M far more than it responds to anything on the exit side: raising it from 0.25 to 0.55 moves the five-year retention rate from 0.68 to 0.64 against a target of 0.54, and more than halves the calibration loss. By comparison, freeing the exit responsiveness across an entire recalibration moved five-year retention by less than one percentage point, and the convex exit rule described above moved it the right way only by breaking the aggregate exit rate.

We do not adopt the higher value, for two reasons. The first is substantive and is the reason ρ_M was set to 0.25 in the first place: an enrollee in the intermediate health state is by construction sicker than a healthy one and should not pass underwriting at close to the healthy rate, whatever it does for the fit. A parameter that is identified by an assumption about who insurers accept should not be re-identified by a moment about how long pools drain. The second is that the response is not clean enough to fit on: it is non-monotone across the values we tried, with $\rho_M = 0.40$ fitting worse than either 0.25 or 0.55 on every retention horizon, which at two seeds per cell is consistent with simulation noise of the same order as the effect. Raising ρ_M also moves the carrier-switching rate further from the UAS figure, trading a validation moment against the targeted ones.

What the exercise does establish is where the limitation lies. The model retains trapped enrollees too long because too few of them can pass underwriting and leave, not because those who would rather drop coverage are insufficiently responsive to price. That is a sharper statement than "the model empties closed blocks too slowly", and it points at the one primitive in this part of the model that is assumed rather than estimated. Consumers who have exited can re-enter if premiums fall relative to their WTP:

$$\Pr(\text{reenter}_i) = \lambda_i \cdot \gamma^{\text{re}} \cdot \max\left(0, \frac{v_i - p_b^{\min}}{v_i}\right). \quad (25)$$

Re-entry requires passing underwriting; the re-entry responsiveness $\gamma^{\text{re}} = 0.50$ is calibrated separately to match the approximately 1.6% annual re-entry rate in the UAS data. The exit and re-entry responsiveness parameters were re-calibrated jointly with the risk-averse WTP: because the risk premium raises WTP, fewer consumers exit and more re-enter at a given premium, so γ^{exit} rises and γ^{re} falls relative to a risk-neutral specification.

G.4 Insurer Pricing

Each block b sets a community-rated annual premium updated each year via exponential a markup over the prior year's claims, as in equation (15). Open plans use markup $m^{\text{nd}} = 1.24$, selected in earlier stages of the grid search and held fixed in the final one, close to the large-company aggregate non-dp markup in the MSI data (1.239), implying a target loss ratio of $1/1.24 = 80.6\%$. Closed dp blocks charge a markup that ramps with years since closure— $m^{\text{nd}} + 0.02$ in the first year after closure, rising to $m^{\text{nd}} + 0.08$ (about 1.34) from the third year on—scaled by a calibrated multiplier on this schedule and in line with the large-company dp-closed aggregate markup in the MSI data (1.294).

G.5 Deadpooling

G.5.1 Block Closure Decision

A block is *deadpooled* when an insurer closes it to new enrollment *and* the same parent opens a fresh block in the same market, matching the empirical definition of Section 2.4; closure alone is an ordinary exit and is not deadpooling. We take the closure decision as exogenous, governed by a hazard rate estimated from the market data. Treating it as exogenous to the individual block is supported rather than assumed: closures are coordinated, subsidiary-wide events whose timing is a corporate decision rather than a response to any one block's pool (Appendix D.1). The annual probability that a parent company initiates a new deadpool in a given year is calibrated to match the observed frequency of block closures in the MSI microdata. Specifically, the deadpooling hazard is modeled as a scaled version of estimates from our calculations from the Mark Farrah data:

$$\Pr(\text{new deadpool in year } t) = h_{\text{MFA}} \times s_h, \quad (26)$$

where h_{MFA} is the estimated baseline (block-age-specific) hazard and $s_h = 1.20$ is a scaling factor. Each open block of a deadpooling parent draws independently against the age-specific hazard; at most one block is closed per state-plan-letter market each year, and if several blocks draw a closure the largest (most-enrolled) is the one closed, reflecting that insurers prefer to maximize the trapped pool.

G.5.2 Closed-Block Premium Dynamics and Lock-in Mechanism

When a block closes, it stops accepting new enrollees. The closed block's risk pool deteriorates through two channels: (i) healthy enrollees exit or switch to cheaper open blocks, and (ii) the remaining pool ages and sickens over time. Closed blocks charge the ramped markup described in Appendix G.4. As healthy consumers depart, average claims rise, triggering further premium increases, which induce more mobile consumers to exit, further deteriorating the pool.

The underwriting step. Because this is the mechanism the whole model turns on, we set it out separately rather than folding it into the switching rule. An enrollee who has decided to move—having shopped, found a cheaper open block, and judged the saving worth taking—must still be accepted by

the receiving insurer. Acceptance is drawn as

$$\Pr(\text{accepted} \mid h, a) = \max\{0, \rho_h - \delta \cdot (a - 65)\}, \quad \rho_H = 0.95, \quad \rho_M = 0.25, \quad \rho_S = 0.05,$$

with $\delta = 0.010$ per year of age above 65. Three features deserve emphasis.

First, the ρ_h are pass rates *at age 65*, not flat rates. Because the age penalty is subtracted rather than scaled, it bites hardest where the base rate is smallest: a sick enrollee's acceptance probability reaches zero at age 70 and a moderate enrollee's at age 90, while a healthy enrollee's is still positive well beyond the top of the age range. Most enrollees in a closed block are therefore facing acceptance probabilities materially below the headline figures, and for the sick the barrier is absolute rather than probabilistic.

Second, the gradient across health states is what generates the selection channel: it is the *difference* between ρ_H and ρ_M , not their level, that determines how much a closing block cream-skims. Flattening the gradient while holding the average fixed removes the mechanism without changing how much friction enrollees face on average.

Third, underwriting binds only outside a guaranteed-issue window. The birthday-rule and continuous-GI counterfactuals of Section 9.2 operate precisely by waiving this step, which is why they act on lock-in without touching the closure decision itself.

The pass rates are calibrated rather than estimated, and Appendix B reports the independent MCBS-and-filings exercise against which they should be judged: on that evidence the average enrollee would be declined by 68.5% of the company $\times$ state markets available to them.

G.6 Simulation Design

G.6.1 Initialization and Timing

The market is initialized with $N_0 = 40,000$ enrollees distributed across all active blocks. To reach a realistic stationary distribution, we discard the first 60 years of the 80-year run, reporting outcomes only over the remaining 20 years. Within each period, premiums carried over from the prior year are in force; then new 65-year-olds enroll; then consumers make switching, exit, and re-entry decisions; claims realize; and insurers update next year's premiums.

G.6.2 Policy Transition Design

Policy counterfactuals are implemented as unanticipated regime changes. The baseline regime runs for the first $T_0 = 20$ years, which also serve as the burn-in and establish the calibrated steady state. At year $T_0 = 20$, the policy regime switches to the counterfactual, and the simulation runs for an additional $T_1 = 60$ years (through year 79) under the new regime. The first T_0 years are the *pre-switch* baseline; they are not the same as the 60 years discarded before reporting, which is the burn-in referred to in Section G.6.

G.6.3 Outcome Aggregation

We report steady-state values averaged over years 60–79, the last 20 years of the 80-year horizon. Every regime, the baseline included, uses this same window; the policy switch occurs at year $T_0 = 20$, so the window opens twenty years after the switch and measures the regime rather than the transition into it. The transitional dynamics are reported separately over the ten years following the switch. Key outcomes include the enrollment-weighted average premium $\bar{P}$, loss ratio $LR = \bar{c}/\bar{P}$, exit and re-entry rates, share out of market, and three nested deadpooling-share measures: the share of enrollees at a deadpooling company (the calibration target, ϕ^{dp}), the share in a block lineage that ever deadpools, and the share currently sitting in a closed deadpooled block. The deadpool premium excess ΔP^{dp} compares premiums in currently-closed deadpooled blocks to the rest of the market.

G.7 Model Parameters

Table A16 reports all structural parameters of the simulation model, their calibrated or externally assigned values, and their sources.

Table A16. Model Parameters: Most Are Externally Calibrated, Three Fit by Grid Search

Parameter	Symbol	Value	Source / Notes
<i>Panel A: Health and Demographics</i>			
Annual claims, healthy / moderate / sick	$\bar{c}_H, \bar{c}_M, \bar{c}_S$	\$846 / \$1,289 / \$2,337	MCBS PAMTALPR ratios 1 : 1.54 : 2.77, NAIC-anchored (\$2,205 avg)
Claims CV, healthy / moderate / sick	CV_H, CV_M, CV_S	2.4 / 2.0 / 1.8	MCBS PAMTALPR; 0–1, 2–3, 4+ chronic conditions
Age claims multiplier, 65 / 75 / 85	$m_c(a)$	1.000 / 1.204 / 1.299	MCBS, normalized at 65
Mortality scale, healthy	μ_0^H	0.015	Relative scale; see age profile below
Mortality scale, moderate	μ_0^M	0.030	SAH RR $\approx 2\times$; Lee et al. (2003)
Mortality scale, sick	μ_0^S	0.080	Assumed; $\approx 5\times$ healthy
Mortality age multipliers	$m_\mu(a)$	see text	Solved to reproduce the SSA life table; $m_\mu(65) = 0.29$, $m_\mu(75) = 0.88$, $m_\mu(85) = 1.45$
$H \rightarrow M$ baseline rate	p_{HM}^0	0.06	MCBS health transition matrix
$M \rightarrow S$ baseline rate	p_{MS}^0	0.09	MCBS health transition matrix
Age health gradient	δ	0.003/yr	Above age 65
<i>Panel B: Willingness to Pay</i>			
WTP horizon	T	20 years	Expected coverage duration from age 65
Risk aversion (CARA)	ρ	1×10^{-4}	Certainty-equivalent WTP loading

continued on next page

Table A16, continued from previous page

Parameter	Symbol	Value	Source / Notes
WTP center, healthy	$\bar{v}_H$	\$1,821	20-yr Markov forward sim, $h_0 = H$ (at $\rho=10^{-4}$)
WTP center, moderate	$\bar{v}_M$	\$2,716	20-yr Markov forward sim, $h_0 = M$
WTP center, sick	$\bar{v}_S$	\$3,580	20-yr Markov forward sim, $h_0 = S$
WTP coefficient of var.	σ_v	0.40	Calibrated to exit/entry moments
<i>Panel C: Price Sensitivity</i>			
λ distribution mean	μ_λ	0.20	Calibrated to exit-rate moment
λ distribution std	σ_λ	0.20	Calibrated
<i>Panel D: Consumer Behavior</i>			
Exit responsiveness	γ^{exit}	0.09	Calibrated (grid search) to the 0.56% UAS annual exit rate
Re-entry responsiveness	γ^{re}	0.50	Calibrated to 1.6% re-entry rate
Shopping probability scale		0.20	Calibrated; applied to λ_i
Savvy attention boost		20.0	Multiplies λ_i in savvy counterfactuals (shop prob capped at 1)
Underwriting pass, healthy	ρ_H	0.95	Industry underwriting tables
Underwriting pass, moderate	ρ_M	0.25	KFF denial rates back-calculated; 18% overall denial
Underwriting pass, sick	ρ_S	0.05	Industry underwriting tables
UW age penalty		0.01/yr	Above age 65; industry tables
Base enrollment, healthy		35%	Calibrated to dp_share moment
Base enrollment, sick		55%	Calibrated to dp_share moment
<i>Panel E: Insurer Pricing and Markup</i>			
Non-DP markup	m^{nd}	1.24	Held fixed (searched in earlier calibration stages); near MSI non-dp markup 1.239
DP-open markup	$m^{\text{dp-open}}$	1.24	Same as non-dp markup
DP-closed markup	$m^{\text{dp-closed}}$	1.28–1.34	Ramps $m^{\text{nd}}+0.02$ to $+0.08$ over yrs since closure (calibrated multiplier); near MSI 1.294
<i>Panel F: Deadpooling Hazard</i>			
P(parent has DP block)	p_{dp}	0.11	Calibrated; MFA estimate of the same quantity is 0.293
Hazard scale factor	s_h	1.20	Fixed at max-trapping (dp_share_closed held out)

Notes: The table’s structure is the point: the health, mortality, claims, and willingness-to-pay primitives (Panels A–B) are set externally from the MCBS, CDC life tables, and UAS rather than fit to market outcomes, so only four parameters (the non-deadpooled markup, the deadpool markup ramp, the hazard scale, and the deadpooling-parent share) are calibrated internally, and the last two are pinned rather than fit. The successor entry price is not listed because it is not identified by any moment; Section 7.2 shows the data contradict the value the model assumes.

G.8 Are the Stylized Dimensions Adequate?

The model abstracts the market to $F = 37$ parents, $S = 5$ states, and $L = 3$ plan tiers, for up to 555 blocks. The tier count is set to three because Plans F, G, and N cover the great majority of lives, but the three modeled plans are representative tiers calibrated to aggregate market moments rather than models of those particular letters: nothing in the calibration is letter-specific, and the model is not asked to reproduce any letter’s benefit design. We justify each choice against the market it stylizes and show that the calibration is insensitive to the two counts that are genuinely free.

Companies, states, and blocks. The real national market contains 239 parent groups, but it is highly concentrated: in 2025 the largest parent holds 34% of covered lives, the top five hold 62%, and the top eight hold 71%. What matters for the deadpooling mechanism is not the long tail of tiny carriers but how many meaningful competitors a beneficiary faces within their own state-and-letter market, and there the median is 11 parents (interquartile range 5 to 30). The model uses 37 parents, above that median but inside the interquartile range. The five states are independent replicas of the no-switching regime the baseline calibrates to (there is no cross-state switching in the model), so the state count governs Monte-Carlo breadth rather than realism.

We report the parent count as a *fit requirement* rather than as a description of the market, because that is what it is. The share of parents that deadpool enters the model only as $\text{round}(N \times \text{share})$, so with a single deadpooling parent the reachable values are multiples of $1/N$, and N is effectively the resolution at which the trapped share can be matched. Reproducing the observed share of lives in a closed deadpooled block requires one deadpooler in roughly thirty-seven. The direction of that requirement is worth noticing: the model needs deadpooling to be done by a *small* carrier, whereas the empirical pattern is the opposite—closures cluster in large national carriers, and closures within a subsidiary are coordinated across letters and states (Appendix D.1). The model matches the aggregate exposure with one small firm where the market generates it with a few large ones. Nothing in the policy counterfactuals turns on which firms deadpool, since the regimes act on the practice rather than on any firm’s identity, but the discrepancy is a limitation of the market structure and not a finding about it.

Robustness to the counts. Average block size, not number of blocks, is what the deadpooling decision’s economics turn on—though the reason needs stating carefully, because a *rate* of return is scale-free and would be unchanged by block size. What size changes is the number of dollars a given markup collects: the excess margin on a trapped pool is a percentage, but the parent weighs it against the largely fixed cost of standing up a successor subsidiary, and only the dollar amount scales with the pool. Holding size fixed is therefore the right invariance to impose. We vary the two free counts while

holding average block size fixed (scaling the simulated population with the block count) and re-run the baseline at the calibrated parameters. Table A17 reports the result: taking the number of parents from 16 to 60, and the number of states from 3 to 10, moves the average premium by less than half a percent, the loss ratio and exit rate essentially not at all, and average insurer profit by under one percent. Two caveats attach to that sweep. It was run under an earlier calibration of the model, which differed from the one in Appendix G.9 in its mortality profile, its mobility targets and the level of claims by health state, so it shows that the aggregate moments are insensitive to the counts under that parameterization; we have not repeated it under the current one. The swept range brackets the 37 parents the calibration selects rather than stopping short of it, so the operative count is inside the range and not an extrapolation from it. The two deadpool-specific micro-moments (the closed-deadpool share and the markup wedge) are noisier because only a handful of closed deadpooled blocks exist in any single run, but they show no systematic trend in either count. The stylized dimensions are adequate: the calibrated market’s behavior is a property of block-level economics, not of how many blocks we simulate.

Table A17. The Calibration Is Insensitive to the Number of Companies and States

Varying	Avg. premium	Avg. claims	Loss ratio	Exit rate	Avg. profit
<i>Companies (S = 5)</i>					
<i>F</i> = 16	\$2,765	\$2,240	0.810	0.0057	\$577
<i>F</i> = 24	\$2,762	\$2,237	0.810	0.0057	\$570
<i>F</i> = 37	\$2,778	\$2,251	0.810	0.0058	\$567
<i>F</i> = 48	\$2,770	\$2,244	0.810	0.0057	\$561
<i>F</i> = 60	\$2,764	\$2,239	0.810	0.0058	\$571
<i>States (F = 37)</i>					
<i>S</i> = 3	\$2,780	\$2,251	0.810	0.0058	\$567
<i>S</i> = 8	\$2,770	\$2,244	0.810	0.0058	\$565
<i>S</i> = 10	\$2,775	\$2,247	0.810	0.0057	\$565

Notes: Baseline simulation at the calibrated parameters, varying the number of parent companies *F* (top block, states fixed at *S* = 5) and the number of states *S* (middle block, companies fixed at *F* = 37), each with *L* = 3 letters. Average block size is held fixed by scaling the simulated population with the block count $F \times S \times L$, so a change in a count is not confounded with a change in block granularity. Each row averages two seeds. The calibration-target row has been dropped: the sweep and the current targets are not stated on the same basis, and the table is a sensitivity across counts rather than a comparison against targets. Aggregate moments and profit are stable to within half a percent across the range; the closed-deadpool share and markup wedge (omitted here) are noisier from small closed-block counts but show no trend in either count.

G.9 Calibration

Most of what the model takes as given is measured. The health and claims primitives come from the MCBS, and mortality from the SSA period life table. Willingness to pay is the certainty equivalent of simulated future claims, not an estimated taste parameter. The underwriting pass rates come from the MCBS eligibility exercise of Appendix B, the two mobility margins (the rate of leaving Medigap and the rate of switching carrier) are taken from our UAS module, and the block-closure hazard is estimated from the NAIC panel. When the model prices harm it uses the *empirically recovered* premium ramp, not its own markup, so the policy magnitudes inherit the credibility of the causal estimates.

We calibrate the model by grid search to match ten moments of the Medigap market, measured on the 2025 NAIC panel for states without guaranteed-issue or birthday-rule protections, the MCBS, the SSA life table, and our UAS module. From the NAIC filings we take the average premium and average claims

Table A18. The Model Matches Prices and the Closed-Block Share Closely, Post-Closure Retention Less Well

Moment	Target	Model	% Error	Weight	Share of misfit
Average annual premium, open blocks	\$2,761	\$2,799	+1.4%	5	<1%
Average annual premium, deadpooled blocks	\$3,373	\$3,393	+0.6%	5	<1%
Average annual claims, open blocks	\$2,164	\$2,256	+4.3%	5	3%
Average annual claims, deadpooled blocks	\$2,551	\$2,584	+1.3%	5	<1%
Share of lives in a closed deadpooled block	0.105	0.108	+3.0%	15	5%
Annual rate of leaving Medigap	0.0056	0.0049	−12.4%	5	28%
Retention, 2 yrs post-closure	0.851	0.876	+3.0%	4	1%
Retention, 3 yrs post-closure	0.724	0.770	+6.4%	4	6%
Retention, 4 yrs post-closure	0.612	0.686	+12.1%	4	21%
Retention, 5 yrs post-closure	0.543	0.630	+15.9%	4	36%
<i>Reported but not targeted</i>					
Annual rate of switching carrier [†]	0.0197	0.1094	+455.1%	—	—
Loss under the full objective: 0.2810; 4 parameters searched.					

Notes: Targets are measured on the 2025 NAIC panel for states without guaranteed-issue or birthday-rule protections, the MCBS, the SSA period life table, and our UAS module; the table is generated directly from the calibration output. Four of the ten targeted moments fit within about three percent (both premiums and both claims) and the share of lives sitting in a closed deadpooled block within four. Two misses are large enough to name. The rate at which enrollees leave Medigap is high by about six percent. The post-closure retention curve is high by five to twenty-five percent, rising with horizon—the model empties closed blocks too slowly. The carrier-switching rate is reported in a separate panel because it is *not* targeted; the model produces several times the annualized UAS figure, and the gap is discussed in Appendix G.9. Every individual route out of a closed block (mortality against the SSA life table, market exit and carrier switching against the UAS) is separately anchored and roughly right on average, so the shortfall is one of *concentration* rather than level: enrollees facing a large post-closure premium increase should leave far faster than the market average and in the model they do not. This biases the model’s lock-in duration, and therefore any harm figure integrating over it, *upward*. The last two columns report a leave-one-out refit: for each targeted moment we drop it from the objective, re-search the whole parameter grid with every lever free, and record whether the calibrated parameters change and what loss the retained moments attain. Only the closed-block share moves them (dropped, the grid selects six deadpooling parents instead of five and a steeper markup ramp) so the calibration rests on that moment and is insensitive to any other single target. The loss column locates the residual misfit: dropping the five-year retention point alone cuts the objective from 0.5421 to 0.2882, so roughly half the total loss is that one moment, yet the argmin is unchanged, because no parameter in the grid reaches it. The switching row has no entries because it is not in the objective and so cannot be left out of it.

of open and of deadpooled-closed blocks separately, the share of enrollees sitting in a closed deadpooled block, and the four-point post-closure retention curve (lives remaining in a closed block two to five years after closure, relative to the first year). From the UAS we take one mobility moment, the annual rate at which enrollees leave Medigap altogether. The annual rate at which enrollees switch carrier while remaining insured is computed and reported alongside the targeted moments but is *not* fit to: the UAS measures it as an “ever switched” indicator, which we annualize over the average exposure of respondents, and that conversion assumes at most one switch per person, a constant hazard over the exposure window, and no attrition of switchers out of the sample. We therefore treat it as a validation check, not a target. Four parameters are searched: the share of parents that deadpool, the scale of the closed-block markup ramp, and the level and the convexity of the exit response to price. The open-block markup is instead held at 1.24, the aggregate premium-to-claims ratio of non-deadpooling, never-closing blocks in unregulated states, so that a fitted markup cannot absorb the loss ratio.

Three primitives are anchored to external data in specific units. First, mortality is anchored directly to the SSA period life table. Enrollees age into worse health states, which already raises their mortality, so scaling each health state’s hazard by the *population* age gradient, which itself rises with age partly because people become sick, would count ageing twice. Instead, the age profile is solved so that population-weighted mortality reproduces the life table at every age, given the model’s own health mix. Second, the two mobility moments are annual *flows*. The UAS asks whether a respondent has ever left

Medigap, a cumulative stock roughly an order of magnitude larger than a yearly rate, so both moments are converted to annual rates using respondents' years since Medicare eligibility. Third, the level of claims by health state is set so that the enrollment-weighted mean reproduces the aggregate claims target under the model's own age distribution; the MCBS ratios across health states are taken as measured. Underwriting pass rates are the MCBS-and-KFF-based values of Appendix B.

Three limitations remain. The first is the firm side. A block prices as a markup over its own recent claims and closes at an estimated hazard rather than solving a dynamic problem, and every parent is identical in size by construction, so the model cannot speak to how insurers would re-optimize entry or pricing strategy in response to a reform, nor to which firms deadpool. Second, the model drains closed blocks more slowly than the data do: every individual route out (mortality, market exit, carrier switching) is separately anchored and roughly right on average, yet a closed block in the model retains too many enrollees at each horizon. Because the *average* mobility rates fit while the *retention curve* does not, the shortfall is one of concentration rather than level: enrollees facing a large post-closure premium increase should leave far faster than the market average, and in the model they do not leave fast enough. At the calibrated parameters the gap widens with horizon, from 3.0% two years after closure to 15.9% at five years. This biases the model's lock-in duration upward, and with it any harm figure that integrates over time: the policy-counterfactual harm in Section 9.2 scores every living beneficiary against the cheapest open premium in their market, and its locked-in and inattentive components come from enrollees paying more than that premium, so a retention curve that is too high at every horizon leaves more enrollees in overpriced closed blocks in the steady state than the panel supports. The harm estimates should therefore be read as upper bounds on the duration dimension.

We have also localized the retention shortfall rather than leaving it general. It is an underwriting margin, not an exit margin: retention responds to the acceptance probability for enrollees in intermediate health far more than to anything governing exit, and a generalized exit rule that concentrates departures among the most heavily repriced was estimated and rejected (Appendix G.3). The model retains trapped enrollees too long because too few of them can pass underwriting and leave, not because those who would rather drop coverage are insufficiently responsive to price. Second, the attention and switching parameters are currently disciplined only indirectly; a new UAS wave scheduled for September 2026 will measure shopping behaviour and carrier switching directly, and both should be re-estimated against it rather than calibrated.

The model fits the targeted moments well on the market's price dimension and less well on its dynamics. Both premiums are within 1.4%, claims within 4.3% for open blocks and 1.3% for deadpooled ones, and the share of lives sitting in a closed deadpooled block (the structural moment the deadpooling story turns on, and the most heavily weighted in the objective) within 3.0%. The rate at which enrollees leave Medigap altogether, anchored to the UAS, is low by 12.4%.

Two misses matter and are not averaged away here. Carrier switching, which is reported but not targeted, runs at several times the annualized UAS figure (0.109 against 0.020) and no calibrated parameter reaches it: the moment has a target but not a lever, which the September 2026 UAS wave is

meant to remedy by measuring shopping and switching directly. The post-closure retention curve is too high, and increasingly so with horizon—3.0% at two years, 15.9% at five. Closed blocks in the model empty too slowly. The miss is diagnostic because every route out of a closed block is separately anchored and individually about right: mortality reproduces the SSA life table, market exit and carrier switching are pinned to the UAS. The average rates fit while the retention curve does not, which locates the problem in the *concentration* of mobility rather than its level—enrollees facing a large post-closure premium increase ought to leave far faster than the market average, and in the model they do not. The consequence runs in a known direction: the model holds trapped enrollees longer than the data do, so its lock-in duration, and any harm figure that integrates over it, is biased *upward*. Mean consumer surplus among enrollees is negative at baseline, about $-\$52$ per enrollee per year. That is a property of the measure rather than a finding about the market. Willingness to pay here is the certainty equivalent of a consumer's *own* future claims—their expected value plus a risk premium—so it contains no value of coverage beyond smoothing that risk, while the premium paid is community-rated and carries the open-block markup of 1.24. Healthy enrollees therefore pay well above their own expected claims by construction, a mean willingness to pay of $\$1,822$ against a mean premium of $\$2,613$, and they are the majority; the pooled mean is negative even though sick enrollees hold a surplus of $\$498$. Levels of this measure are accordingly not interpretable as welfare, and only differences across policy regimes are.

The grid searches four parameters. At the selected values 4 of the model's 37 parent companies deadpool (11%), the closed-block markup ramp has a multiplier of 1.05 (a multiplier on the schedule of excess markups a closed block charges as it ages), and exit responds to price with responsiveness 0.09 and convexity 0.50. After the final round of the search none of the four lies on the edge of its grid. Two further parameters are held fixed rather than searched: the open-block markup, at 1.24, and the closure-hazard scale, at 1.20. Both were searched in earlier stages of the calibration and settled in the interior of their ranges.

A leave-one-out exercise on an earlier version of this calibration, which searched the parent share, the ramp and the level of the exit response, found that the share of lives in a closed deadpooled block was the one moment doing any identifying work: dropping it moved the grid to a different number of deadpooling parents and a steeper markup ramp, whereas dropping any other moment left the selected parameters where they were. We have not yet repeated the exercise on the current calibration. The willingness-to-pay and price-sensitivity distributions, the entry friction by health state, and the health-transition and mortality parameters are calibrated externally from MCBS and UAS data and held fixed throughout.

G.10 Alternative Measures of Harm in the Model

Section 9.2 scores every beneficiary against the baseline's cheapest open premium, so its harm measure counts everything that keeps people from the best open price, whether or not deadpooling causes it. This appendix defines the measure of harm from deadpooling alone shown in Figure 13 and reports an alternative split of overpayment between lock-in and inattention, from the same simulation run.

Table A19 collects both.

Harm from deadpooling alone. For each person in a closed deadpooling block, we set a counterfactual price equal to what they pay now, scaled down to the markup charged by every other block in the market that year (premium revenue over the previous year's claims). Harm has three parts. Enrollees still in the block pay their premium minus that counterfactual price. Enrollees who left a closed deadpooling block for the outside option, but value coverage above their counterfactual price, lose their willingness to pay minus that price in each year their former block is still open to them as a closed deadpooling block. Each switch out of a closed deadpooling block costs the \$30 search cost used in Section 7.3. Because the counterfactual moves with each block's own claims, an older or sicker pool raises both prices together and does not register as harm. The measure is zero by construction when no block deadpools: under a ban it is 0% of today's level. A birthday rule leaves 74% of it, attentive consumers alone leave 52%, and attention combined with a birthday rule leaves 4%. Automatic switching leaves 3%. Capping closed blocks' premiums at the open-market rate gives a negative figure, -167%, because the cap holds closed blocks' markup below that of open blocks; we do not set it to zero.

Lock-in and inattention by underwriting probability. Figure 14 classifies each overpayer by what they did that year: locked in if they shopped and were rejected by underwriting, inattentive if they did not shop, and a third group who shopped but still pay more than the cheapest open premium. An earlier version of the paper instead split the same overpayment by each enrollee's probability of passing underwriting, counting a share one minus that probability as lock-in and the rest as inattention. Total harm is the same under both; only the split differs. Under the probability weights, lock-in is 52% of today's total harm and inattention 28%, against 14% and 57% when each person is classified by behavior. The probability weights count as locked in everyone who would fail underwriting if they tried, whether or not they try; most overpayers do not shop in a given year, so far fewer are actually turned away. The last two columns of Table A19 report the probability-weighted split for every regime.

Table A19. Alternative Measures of Harm Under Each Reform

Regime	Deadpooling-only harm, % of today's	All harm, % of today's total					
		Total	By behavior			Probability weights	
			Locked in	Inattentive	Shopped	Locked in	Inattentive
Today (baseline)	100	100	14	57	9	52	28
Birthday rule	74	95	0	56	20	0	75
Ban deadpooling	0	100	11	46	8	42	24
Price capped at open-market rate	-167	84	12	47	7	42	24
GHHW long-term contract	23	72	3	64	1	40	28
Automatic switching ($\geq 10\%$ savings)	3	47	0	0	41	25	16
Attentive consumers	52	89	36	13	17	44	22
Attentive + birthday rule	4	17	0	15	21	0	36

Notes: Calibrated model, steady state (years 60–79), four seeds, the same run as Figure 14. First column: harm from deadpooling alone (stayers' premium above the open-market markup, the lost surplus of those who left a closed deadpooling block but would stay at that markup, and switching costs out of closed deadpooling blocks), as a percent of today's. Other columns: total harm and its overpayment components as percents of today's total harm. *By behavior* (as in Figure 14): locked in = shopped and rejected by underwriting; inattentive = did not shop; shopped = shopped but still paying more than the cheapest open premium. *Probability weights*: overpayment split by the underwriting pass probability instead. Priced-out and price-level components are not shown, so the components do not sum to the total.

Harm per affected person. Figure 14 divides each source by every living beneficiary. Table A20 divides it instead by the people it actually hits, with the share of beneficiaries affected in parentheses; locked-in, inattentive and shopped heads count the people in each group that year. Under the baseline, a sick person priced out of Medigap loses \$2,712 a year (7.2% of beneficiaries), and a locked-in enrollee overpays \$2,059 (12.5%). Across regimes, the reforms mostly change *how many* people are hurt rather than how much each affected person loses. The shares in parentheses are of all living beneficiaries, not of each health group.

Two patterns in the attentive-consumer rows follow from how the model represents attention. First, more sick enrollees end up priced out. In the model, greater attention raises not only how often consumers shop but also how readily they drop coverage when their premium exceeds their willingness to pay, and a sick enrollee who drops coverage almost never passes the underwriting needed to return. Pool deterioration adds to this only in closed deadpooling blocks, where the healthy leave faster; open premiums fall. Second, more enrollees are counted as locked in. A locked-in enrollee is one who shopped that year and was rejected by underwriting, so when far more consumers shop, more of them are rejected, even with the underwriting rules unchanged.

Table A20. Harm per Affected Person, by Source and Regime

Regime	Priced out			Locked in	Inattentive	Shopped
	Sick	Medium	Healthy			
<i>Today (baseline)</i>	\$2,712 (7.2%)	\$1,876 (6.3%)	\$688 (6.9%)	\$2,059 (12.5%)	\$2,070 (50.0%)	\$2,024 (8.0%)
Birthday rule	\$2,741 (7.1%)	\$1,885 (6.4%)	\$712 (7.2%)	—	\$2,039 (49.9%)	\$1,893 (18.8%)
Ban deadpooling	\$2,721 (7.8%)	\$1,881 (6.9%)	\$712 (7.7%)	\$1,730 (12.1%)	\$1,740 (48.7%)	\$1,764 (8.2%)
Automatic switching ($\geq 10\%$)	\$2,751 (5.1%)	\$1,886 (4.4%)	\$682 (4.9%)	\$727 (0.0%)	\$684 (0.4%)	\$1,471 (51.2%)
Attentive consumers	\$2,517 (13.4%)	\$1,672 (8.4%)	\$897 (9.2%)	\$2,324 (28.3%)	\$2,242 (10.4%)	\$2,432 (12.8%)
Attentive + birthday rule	\$2,507 (6.7%)	\$1,678 (3.5%)	\$946 (2.3%)	—	\$2,649 (10.1%)	\$2,634 (14.7%)

Notes: Harm per affected person per year in the calibrated model, January 2025 dollars, same run and yardstick as Figure 14; in parentheses, the share of living beneficiaries affected. *Priced out:* out of Medigap although they value it above the baseline's cheapest open premium, by current health. *Locked in, Inattentive* and *Shopped:* paying more than the cheapest open plan, having that year shopped and been rejected by underwriting, not shopped, or shopped without moving to the cheapest open premium. —: no one affected.

Which policy would the model's population choose? We group everyone alive in the last simulated year into types by health, attention and age band; each type votes for the policy that gives people of that type the highest average consumer surplus in that year (counting those out of Medigap at zero), and votes are weighted by the number of living people of each type. Among the regimes in Figure 14, Attentive + birthday rule wins with 81% of the vote and beats every alternative head to head. Attention, however, is not something a regulator can legislate; restricting the menu to rules that can be imposed directly, Automatic switching ($\geq 10\%$ savings) beats every alternative, with at least 100% of the vote in each head-to-head comparison. A ban on deadpooling would not command a majority against the status quo: 46% of the population is better off under it. Because the branches of the simulation consume random numbers differently after the policy switch, the same person is not comparable across regimes, which is why voters are types rather than individuals.