Munich Personal RePEc Archive

Demography in a K-Shaped Nation: How Population Change Will Reshape American Economic, Political, and Social Systems, 2026-2050

LANE, Edward

7 August 2026

Online at https://mpra.ub.uni-muenchen.de/130387/
MPRA Paper No. 130387, posted 03 Sep 2026 11:20 UTC

Demography in a K-Shaped Nation: How Population Change Will Reshape American Economic, Political, and Social Systems, 2026-2050

Edward Lane, ASA, CFP®

Working draft, August 2026

Abstract

America's demographic challenge is not simply that the population is aging or that population growth is slowing. It is that the nation's principal economic, political, and social institutions were built during eight decades of sustained demographic expansion and continue to operate under assumptions of growing workforces, expanding tax bases, rising student enrollments, and robust family support networks. This paper argues that the central issue through 2050 is therefore one of **institutional adaptation** rather than demographic change alone.

The paper develops an integrated analytical framework that examines how demographic deceleration propagates through federal entitlement programs, labor markets, education, political representation, state and local government finance, housing, healthcare, and household caregiving. It introduces the concept of a **K-shaped nation** to explain why similar demographic pressures produce increasingly unequal outcomes across regions and communities. Communities with strong fiscal resources, diversified economies, and resilient institutions are generally better positioned to adapt, while those with weaker institutional capacity face self-reinforcing cycles of population loss, labor shortages, service contraction, and fiscal stress.

A central finding is that the burden of demographic adjustment progressively shifts downward through American society. Federal fiscal pressures become state budget challenges. State constraints become municipal service reductions. Municipal limitations increasingly become household responsibilities. At each stage, the institution receiving the burden typically possesses fewer financial, institutional, and human resources than the institution passing it along.

The paper distinguishes between financial constraints and real resource constraints. While fiscal policy can redistribute claims on future output, it cannot by itself create caregivers, healthcare capacity, housing, or functioning local institutions. Immigration can expand the working-age population within a relatively short period, whereas fertility policies require decades to affect labor supply. Artificial intelligence may increase productivity in many sectors but is unlikely to eliminate shortages in occupations requiring sustained physical presence, including caregiving, nursing, construction, and other essential services.

The paper concludes that America's demographic future is unusually predictable because the cohorts that will shape the next two decades have largely already been born. The principal policy challenge is therefore not forecasting demographic change, but adapting institutions built for an era of continuous growth to one of demographic deceleration and, in many places, demographic contraction. Whether demographic change becomes a

manageable transition or an increasingly unequal redistribution of economic and social burdens will depend largely on the nation's capacity for institutional adaptation.

Keywords: Demographic Change; Population Aging; Institutional Adaptation; K-Shaped Nation; Institutional Capacity; Social Security; Medicare; Medicaid; Long-Term Care; Labor Force; Immigration; Artificial Intelligence; State and Local Government Finance; Regional Divergence.

JEL Classification: J11, J14, J21, J61, H55, H75, R23, I38.

Synopsis

America's demographic challenge is often described as a problem of aging, declining fertility, or slowing population growth. This paper argues that those descriptions identify the symptoms rather than the underlying issue. The deeper challenge is that many of the institutions on which Americans depend—including schools, colleges, retirement systems, healthcare programs, state and local governments, housing markets, and family caregiving networks—were built during nearly eight decades of sustained demographic expansion. As population growth slows and, in many communities, begins to reverse, those institutions must operate under conditions for which they were never designed.

The paper examines how demographic change will reshape American economic, political, and social systems through 2050. Rather than treating aging, labor markets, healthcare, education, housing, and public finance as separate policy issues, it develops a unified framework that follows demographic change as it moves through interconnected institutions. The analysis argues that demographic change is best understood as a problem of institutional adaptation.

A central contribution of the paper is the concept of a K-shaped nation. While demographic aging will affect virtually every region of the country, its consequences will not be evenly distributed. Communities with strong tax bases, diversified economies, robust healthcare systems, higher educational attainment, and greater household wealth possess buffers that enable them to adapt successfully. Communities with fewer institutional resources face a different trajectory. Population loss, labor shortages, shrinking tax bases, declining public services, and weaker economic opportunities reinforce one another, producing increasingly divergent outcomes despite similar demographic pressures.

The paper argues that demographic adjustment also follows a distinct institutional path. Fiscal pressures that first appear at the federal level increasingly shift to states. State governments, constrained by balanced-budget requirements, pass many of those pressures to municipalities. Local governments facing declining tax bases reduce services or defer investment. Ultimately, many responsibilities are transferred to households, where family members provide care, financial support, or unpaid labor that formal institutions can no longer supply. At each step, the recipient generally possesses fewer resources than the institution passing the burden along.

Particular attention is devoted to the nation's major social insurance programs. The paper distinguishes the fundamentally different demographic challenges confronting Social Security, Medicare, and Medicaid and explains why each responds differently to aging and slower labor-force growth. It argues that the long-term constraint extends beyond public

finance to the availability of workers capable of providing care. Financing reforms may redistribute costs, but they cannot create caregivers, healthcare capacity, accessible housing, or functioning local institutions where those resources do not exist.

The analysis also considers the principal factors capable of influencing future demographic outcomes. Immigration is identified as the only policy capable of materially increasing the working-age population within the next decade, while fertility policies require a generation to influence labor supply. Artificial intelligence may significantly improve productivity in many sectors, but it is unlikely to eliminate shortages in occupations requiring sustained human presence, including caregiving, nursing, construction, and other hands-on work.

The paper further considers how improvements in healthspan, regional migration, and climate-related geographic shifts may influence the pace and distribution of demographic adjustment without fundamentally altering the underlying demographic arithmetic.

Rather than presenting demographic change as an unavoidable crisis, the paper argues that the country's future remains highly responsive to institutional choices. America's demographic trajectory is unusually predictable because most of the people who will shape the next two decades are already alive. The principal uncertainty lies not in forecasting population trends, but in determining whether institutions built for continuous expansion can successfully adapt to an era of slower growth and demographic contraction.

The paper concludes that successful adaptation will depend less on financial accounting than on expanding the nation's real productive capacity: strengthening the caregiving workforce, increasing housing supply, improving labor-force participation, investing in education and workforce development, and enhancing the resilience of state, local, and community institutions.

The central question is therefore no longer whether demographic adjustment will occur, but who ultimately bears its costs if institutions fail to adapt. The answer will determine whether America's demographic transition becomes a period of orderly adjustment or an increasingly unequal redistribution of burdens across governments, communities, and households.

Contents

1. Introduction: The Most Predictable Crisis in American Life 11

1.1 A Note on the Baseline 16

2.1 The Boomers Are Already Here 18

2.2 Why Birth Rates Aren't Coming Back 20

2.3 When Deaths Start Outnumbering Births 22

2.4 Immigration: The Only Lever That Moves Fast 23

2.5 Where People Are Going, and What They Leave Behind 27

2.6 Who the Country Is Becoming 28

2.7 These Problems Feed Each Other 29

2.8 Where My Framework Doesn't Work 32

2.9 If the Door Stays Closed 32

3. Three Programs, Three Different Failures 34

3.1 What a Trust Fund Date Actually Means 35

3.2 Social Security: Six Years to Decide 37

3.3 Medicare: Everyone Is Watching the Wrong Number 41

3.4 Medicaid: The Program That Gets Hit Hardest 43

3.5 You Can't Appropriate Caregivers 45

3.6 What Aging Does to Asset Prices 46

3.7 Where That Leaves Us 47

4. Two Shocks, Moving in Opposite Directions 49

4.1 The Workforce Stops Growing 49

4.2 Where the Jobs Actually Are 51

4.3 What We Actually Know About AI and Jobs 52

4.4 What Happens When You Delete the Bottom Rung 53

4.5 Why the Two Shocks Don't Cancel Out 54

4.6 Scarcity Changes How Things Get Made 56

4.7 Who Pays, and Where 57

5. Education: The Pipeline Empties on Schedule 57

5.1 The Schools Are Already Emptying 58

5.2 Fewer Students Won't Fix the Teacher Shortage 60

5.3 The Enrollment Cliff Is Here 61

5.4 Closing the Schools We Are About to Need 62

5.5 Where That Leaves Us 63
6. How Population Turns Into Political Power 64
6.1 The Reapportionment Machine 64
6.2 What Callais Changed 66
6.3 Having the Numbers Isn't Having the Votes 67
6.4 Everything Lands at Once 68
6.5 The Political Paradox 69
6.6 Where That Leaves Us 71
7. The Governments That Cannot Print Money 71
7.1 What Happens to the Tax Base 72
7.2 The Costs That Do Not Shrink 73
7.3 Pensions and Other Promises 73
7.4 What Shows Up in the Credit Markets 74
7.5 Why Some States Can Absorb This and Others Cannot 74
7.6 What Managing Decline Actually Looks Like 75
8. What This Looks Like at the Kitchen Table 77
8.1 The Wrong Houses in the Wrong Places 77
8.2 Most Care Is Unpaid, and Nobody Counts It 78
8.3 The Insurance Market That Never Formed 80
8.4 Who Takes Care of You When Nobody Is Left 81
8.5 The Wealth Transfer and Why It Widens the K 82
8.6 Where That Leaves Us 84
9. What Would Actually Help 84
9.1 How I Am Judging These 84
9.2 The Zero-Sum Trap 87
9.3 Immigration 89
9.4 The Care Economy 89
9.5 Housing 91
9.6 Social Security and Medicare 92
9.7 Four Things That Will Not Work 92
9.8 Where the Positions Actually Overlap 93
9.9 What Is Doable Is Not What Works 96
10. The Consolidated Timeline 97

11. Conclusion 100
Appendix A. Implications for Institutional Investors 103
A.1 A Demographic Credit Screen 103
A.2 Asset Implications 104
References 107

Executive Summary

America is getting older and may eventually become smaller. Meanwhile, its institutions were built for continuous population growth and do not know how to adapt when that growth stops. That is the problem this paper addresses.

School revenue falls with enrollment, while buildings, transportation systems, and administrative costs remain. Colleges depend on a replenishing supply of students and tuition. Social Security relies on payroll tax financing as the worker-to-beneficiary ratio falls. State and local governments must maintain infrastructure and legacy obligations against shrinking tax bases. These systems lose revenue, workers, students, and taxpayers faster than they shed costs.

The demographic change is already underway. Births have run near 1.6 per woman for a decade. Around 2030, deaths begin to exceed births, labor force growth slows sharply, and immigration becomes the only source of national population growth. Under the Congressional Budget Office's path, the country adds about 13 million people over the next 25 years. If current immigration policy holds, it adds close to none. In the 25 years after 1980, it added nearly 70 million.

Total population decline is only one possible outcome. The more immediate facts are already fixed. Fewer children are entering schools. Smaller cohorts are approaching college and working age. The number of beneficiaries is rising faster than the number of workers supporting them. Many counties are already losing population. Institutions can contract even while total national population continues to grow slowly.

Immigration determines how quickly pressure builds, as it adds working-age people within a few years, unlike fertility and mortality, which take decades. The difference between low- and high-growth scenarios is about 25 million workers by 2050, roughly the size of the combined labor forces of Texas and Florida. Factors like age at arrival, skills, family makeup, and settlement location shape the impact on labor markets, schools, and politics.

The consequences will be K-shaped. Affluent areas also experience aging, fertility decline, and enrollment loss, but they possess universities, diversified hospital systems, asset-based tax bases, fiscal reserves, access to capital, and continuing in-migration. Those resources give them choices.

Less affluent communities often have one hospital, one college, one school system, and a tax base tied closely to wages and population. When one institution contracts, few

alternatives remain. Hospital closures make a community less attractive to working-age families. Their departure weakens the tax base and reduces school enrollment. School consolidation and declining services then encourage more families to leave. The initial demographic shock becomes a continuing cycle of contraction. The K-shaped result describes which places can interrupt that cycle.

The financial impact also moves downward. While the federal government can change its financing rules, states must balance their budgets and borrow in a currency they cannot issue. Municipalities must maintain water systems, roads, emergency services, and public buildings that cannot be resized quickly. Households have the least ability to pool risk.

Medicaid shows how the process works. A reduction in federal spending becomes a state budget problem. The state responds through reimbursement, eligibility, or service changes. Providers reduce capacity. A family member supplies the missing care without pay. At each step, the recipient has less flexibility than the institution passing the burden along.

Care is where the demographic problem becomes a physical constraint. Direct-care occupations are projected to generate approximately 9.7 million openings between 2024 and 2034. Most are replacement openings caused by turnover and labor force exits, but every opening must be filled to maintain existing capacity. Median annual earnings remain under $26,000, and immigrants supply roughly a quarter of the workforce.

Unpaid relatives absorb the care the formal system does not provide, and the scale of that transfer is easy to miss. It runs to roughly $1.01 trillion a year in replacement value, more than all federal, state, and local Medicaid spending combined. That capacity will decline as families become smaller, relatives live farther apart, and more older adults reach advanced age without adequate support networks. Artificial intelligence may raise national productivity, but it does not put a caregiver in the room. The work most urgently needed requires physical presence.

Policy should therefore expand the country’s ability to provide what people need. Financial reform remains necessary for Social Security and Medicare, but financing alone cannot produce caregivers, housing, clinical capacity, or functioning local institutions.

Immigration can increase the working-age population. Higher direct-care compensation can improve recruitment and retention if Medicaid reimbursement rises with it. Training support can preserve nursing and technical programs in regions where enrollment-based financing would otherwise close them. Housing reform can provide family-sized housing for younger households and accessible housing for older residents. Caregiver

compensation and Social Security credits can reduce the lifetime penalty imposed on relatives who leave paid employment to provide care.

Some communities will not regain their previous populations. Their governments will need to maintain acceptable services with fewer residents through consolidation, regional service delivery, infrastructure reduction, and investment in viable service centers.

The time available for gradual adjustment is narrowing. The 2030 census will be followed by reapportionment and redistricting. The resulting congressional map and electoral votes take effect in the 2032 presidential election. Social Security's retirement trust fund is projected to reach depletion in the fourth quarter of 2032, when continuing revenue would cover 78 percent of scheduled benefits under current law. Medicare Hospital Insurance follows in the second quarter of 2033.

Congress can change the law, and the dates may move. Delay still matters. Changes adopted early can be phased in across cohorts. Changes adopted near depletion fall more heavily on current workers and beneficiaries.

Nothing in this analysis depends on a surprise. The cohorts are born, the dates are known, and the arithmetic has been visible for years. America cannot choose whether demographic change arrives. It can choose whether institutions built for growth adapt deliberately, or whether the adjustment is passed, step by step, to the governments, communities, and families least able to bear it.

1. Introduction: The Most Predictable Crisis in American Life

Everyone who will turn 65 before 2050 is already alive and has already celebrated a 41st birthday. Everyone who will enter the labor force before 2044 has already been born. Demography is the rare branch of economic forecasting in which the long run is more certain than the short run. Recessions arrive unannounced, but the age structure of the population in 2040 is largely arithmetic performed on the number of people who already exist. That certainty is what makes the present moment unusual. The United States can see its demographic future with rare clarity, and that future differs from anything in the national experience.

So what is driving this? Three forces. The first is aging: by 2030 every member of the baby boom generation is 65 or older, and within a decade older adults outnumber children for the first time in American history. The second is fertility. The total fertility rate peaked at 2.12 births per woman in 2007, fell to approximately 1.6 by the mid-2010s, and has stayed there through expansion, pandemic, recovery, and inflation shock alike; CBO now carries 1.6 through mid-century (CBO 2026).

The third is immigration, transformed from a steady tailwind into the most volatile variable in American population arithmetic. Net migration collapsed in 2025 as a result of enforcement policy rather than economic conditions, and by an amount the published estimates do not agree on. Section 2.4 sets out the dispute. Whether that proves a two-year episode or a durable regime change is the largest single source of uncertainty in the population and labor force baseline of this paper.

These shifts arrive alongside rapid advances in artificial intelligence, which will restructure labor demand as demography restructures labor supply, and will do so asymmetrically: upper-branch workers hold the assets that capture automation gains, while lower-branch workers face displacement without offsetting capital income. Section 4 treats the collision directly.

A word on terminology, since I use “contraction” throughout and the word can mislead. Total national population is not projected to fall before 2050 under the baseline; it reaches roughly 362 million. What contracts is more specific: natural increase ends around 2030, cohorts of children and young adults shrink on a schedule already fixed by births, labor force growth slows sharply and stops altogether under low immigration, the worker-to-beneficiary ratio declines continuously, and a growing share of counties loses population outright. Those are the quantities growth-built institutions depend on. A country can add residents while every institution calibrated to rising numbers of children, workers, and taxpayers experiences decline. Outright national decline appears only under the low-immigration scenario of Section 2.4.

And none of this is far off. Social Security’s Old-Age and Survivors Insurance trust fund is forecast to reach depletion in the fourth quarter of 2032, six years from this writing, at which point continuing revenue covers 78 percent of scheduled benefits (Social Security Trustees 2026). The Medicare Hospital Insurance fund follows in mid-2033. The higher education enrollment cliff, the delayed echo of the post-2007 birth decline, is arriving on campuses now. The care economy is already the fastest-growing source of employment demand even as the population available to staff it grows more slowly each year. These are

projections about the current decade, with compounding effects that run through mid-century.

So what is the actual problem? Not aging itself. Aging is manageable, and other rich countries have managed worse age structures than ours. **The problem is that nearly every American institution examined here was designed and financed during continuous population growth, and each has a growth assumption embedded in its mechanics.** School funding is per-pupil, so falling enrollment cuts revenue faster than cost. Social Security is pay-as-you-go. Apportionment is zero-sum against 435 seats, so a state can grow and still lose power. Colleges assume a replenishing pipeline of eighteen-year-olds.

None of these reliably sheds cost as fast as it sheds revenue when the population stops expanding. Some adapt, and a few do so well, but the adjustment is asymmetric and it is not free. So the question through 2050 is not how to slow aging, which we cannot do. It is whether institutions built for growth can be adapted to contraction, and who gets stuck with the cost while we argue about it.

This diagnosis is not idiosyncratic. Two recent literatures reach it independently. Sutradhar, Spearing, and Derrible (2024), projecting American city populations to 2100, find that close to half of the roughly 30,000 U.S. cities face some population decline, affecting 27 to 44 percent of populated area, and conclude that resource and service problems will persist unless planning moves away from growth-based assumptions.

The OECD's study of depopulating regions across member countries reaches a parallel conclusion from the fiscal side, documenting a vicious circle in which eroding tax bases and lost economies of scale degrade services, which accelerates the depopulation that started the cycle (OECD 2025b). Section 7 develops the American subnational version of that argument.

That question has no single national answer, which is the paper's second claim. By the K-shaped economy I mean a structural divergence between an upper branch of generally metropolitan, asset-owning, highly educated households whose fortunes increasingly track AI-exposed equity markets and professional labor demand, and a lower branch of wage-dependent, renter-heavy households concentrated in non-metropolitan and legacy-manufacturing counties, where real consumption capacity has deteriorated despite aggregate growth (Lane 2026a).

Table 1 summarizes the two branches along the dimensions that matter here; Lane (2026a) operationalizes the division at the county level using income, educational attainment,

migration flows, and housing cost-to-income ratios. The framework's claim in this paper is specific, and Section 2.8 states its limits: the K predicts capacity to absorb demographic shocks more reliably than it predicts where those shocks land. Aging is a manageable fiscal problem in a county gaining educated young in-migrants and a compounding catastrophe in a county that has lost them, not because the aging differs but because the absorptive capacity does.

Table 1. Two dimensions of the K: household position and place-level capacity

Dimension	Upper branch	Lower branch
Net migration	In-migration of young, educated workers	Out-migration of the same workers
Age structure	Younger; continuously refreshed	Older; aging accelerated by departures
Income base	Assets plus wages	Wages only
Housing tenure	Owner-heavy; appreciation builds wealth	Renter-heavy; cost burden delays family formation
Fiscal base	Expanding tax base	Shrinking tax base; rising transfer dependence
AI exposure	Exposure to equity upside and professional labor demand	Displacement risk without offsetting asset income
Health outcomes	Life expectancy at the top of the national distribution; screening-stage diagnosis	Life expectancy comparable to middle-income countries; late-stage diagnosis; obstetric and behavioral health deserts
Absorptive capacity	Universities, hospitals, capital markets, and in-migration available to cushion shocks	Few institutional or fiscal buffers; shocks pass straight through to households

Source: author, condensed from the county-level framework in Lane (2026a). The two dimensions are correlated but distinct, and the paper treats them separately where they diverge. A renter in a strong metropolitan economy occupies a weak household position in a place with high capacity; an outright homeowner in a declining county holds an illiquid asset in a place with low capacity. Demographic consequences are most severe where a vulnerable household sits in a place that also lacks capacity.

This is the thesis: demographic change in America will not produce a single national outcome but a widening divergence between upper-branch and lower-branch America. One branch will be labor-constrained but institutionally cushioned. The other will be fiscally fragile and population-declining. Immigration policy is the decisive swing variable, and the division runs within generations and within metros as much as between regions.

One framing choice needs to be stated up front, because it runs through Section 3 and beyond. Two accounts of the fiscal problem are available. The conventional public finance account holds that population aging increases promised benefits relative to dedicated revenue, requiring some combination of higher taxes, lower benefits, borrowing, or offsetting spending cuts. The functional finance account, running from Lerner through Wray, Kelton, and Tcherneva, holds that a currency-issuing government faces no inherent financing constraint, though appropriations, debt limits, and other statutory rules remain operationally binding until changed.

The choice among those options therefore allocates claims on output rather than answering the question of affordability. What binds instead is whether the economy can produce enough goods, services, and labor to satisfy those claims without unacceptable inflation (Lerner 1943; Wray 2015; Kelton 2020).

This paper adopts the second framework without denying the distributional and inflation consequences the first emphasizes. Functional finance does not dissolve the fiscal problem. It relocates it, from the adequacy of a ledger to the adequacy of productive capacity, and it insists that the distributional question, who bears the adjustment, be answered explicitly rather than left implicitly in an accounting trigger.

The paper proceeds as follows. The order is deliberate. The paper moves from national demographic arithmetic to federal institutions, then through labor, education, and politics, then down through state and local government to the household where the residual costs ultimately land. Sections 3 to 5 cover the fiscal and real-economy systems, Section 6 the political system, and Section 7 state and local government. Section 8 follows the adjustment into the household, and Sections 9 and 10 take up policy and the consolidated timeline.

By section:

- Section 2 constructs the demographic baseline for 2026-2050, establishing which combinations of fertility, mortality, immigration, and internal sorting are arithmetically possible.
- Section 3 examines Social Security, Medicare, and Medicaid.
- Section 4 turns to employment, labor markets, and the AI interaction.
- Section 5 addresses education from the K-12 contraction through the higher education consolidation wave.
- Section 6 analyzes reapportionment and the electorate.
- Section 7 examines state and local government, which faces the same demographic pressure without the federal government's monetary latitude.
- Section 8 considers housing, household care, the shrinking family networks on which care depends, and the intergenerational transfer of wealth.
- Section 9 evaluates policy responses.
- Section 10 consolidates the dated milestones.

1.1 A Note on the Baseline

The word baseline needs care, because it sounds more neutral than it is. CBO projects that net immigration recovers after 2025, and the basis for that projection is worth stating precisely. It is not a reading of current law. CBO attributes the 2025 collapse to a June 2024 executive order and to subsequent administrative actions, then projects the affected category to return to an amount consistent with its long-run historical average (CBO 2026). The recovery is a mean-reversion assumption: that the present policy posture is a departure from the historical norm, and that departures decay. That may prove right. But it is a judgment about how long an executive policy will hold, and no statute compels the outcome in either direction.

I have handled that by carrying three immigration paths rather than one. Table 2 sets them out and Table 5 shows what they do to the labor force. Where a number depends materially on which path holds, I say so. Where I report a figure without qualification it follows the CBO path, and the reader should treat it as one of three rather than as the expected case.

The low path is not a worst case or a pessimism exercise. It is what happens if current administrative policy persists, which requires no action by anyone.

One constraint on this paper should be stated rather than concealed. The Trustees Reports, the Bureau of Labor Statistics, and the Census Bureau all build their projections on immigration assumptions comparable to CBO's. Reporting figures from those sources while asserting a different immigration path would produce numbers that do not belong together. I therefore report CBO-consistent figures throughout, not because I judge that path most likely, but because mixing baselines across sources yields incoherence. Section 3.2 estimates separately what the low path would do to Social Security's finances, using the Trustees' own published sensitivity.

A note on method. Baseline projections follow the intermediate assumptions of the 2026 Social Security and Medicare Trustees Reports, CBO's Demographic Outlook published in January 2026, and the Census Bureau's national projections, supplemented by provisional vital statistics through 2025. Where CBO's 2026 outlook supersedes assumptions in its 2025 update, particularly on fertility and immigration, the 2026 figures govern. Where sources disagree, the disagreement is reported rather than averaged.

Throughout, claims are sorted into three classes, and the milestone tables label each entry accordingly. Near-certain claims are arithmetic performed on cohorts that already exist: the number of 18-year-olds in 2040 is fixed by births that have already occurred, subject only to mortality and migration at the margin. Projection-dependent claims are model outputs sensitive to assumptions that reasonable analysts dispute, such as the year national deaths first exceed births.

Policy-contingent claims describe outcomes determined by legislative choices not yet made, including every trust fund depletion consequence in Section 3, since each assumes Congress does nothing. Keeping these separate is the main discipline this paper imposes on itself. Dated estimates carry a standing caution regardless of class: depletion dates and crossover points have moved one to three years between annual projection vintages, and the 2025-2026 immigration shock moved several of them in a single year.

2. What We Already Know Is Coming

Start with what is not in dispute. The United States entered 2026 with roughly 350 million residents on CBO's projection basis, which uses a broader definition than the Census

Bureau's resident-population estimate of about 342 million in mid-2025 and is not interchangeable with it and the slowest peacetime population growth in its recorded history apart from the pandemic year: 0.5 percent, or roughly 1.8 million people, between mid-2024 and mid-2025, with the deceleration driven almost entirely by the collapse of net international migration (U.S. Census Bureau 2026). CBO projects the population will reach only about 364 million by 2056 (CBO 2026). The components of that near-stagnation, and the ways they could break differently, are the subject of this section.

2.1 The Boomers Are Already Here

The baby boom generation, born 1946-1964, numbered roughly 76 million at its peak and still numbers well above 60 million. Its youngest members turn 65 in 2029; by 2030 the entire generation is of retirement age. The share of the population 65 and older, roughly 18 percent in 2026, rises to approximately 23 percent by mid-century.

More consequential than the 65+ share is the composition within it. The population 85 and older, the group that drives long-term care demand, disability rates, and Medicaid spending, more than doubles between the mid-2020s and 2050 as the boomer wave moves through its ninth decade.

The heaviest caregiving burden therefore does not arrive with boomer retirement in the 2020s. It arrives in the late 2030s and 2040s, when the largest cohorts reach 85 and older, an age group in which roughly six in ten adults report some limitation in activities of daily living. Section 3 speaks to what this means for Medicaid, which is, contrary to its public image, the nation's primary payer of long-term care.

One large uncertainty sits inside these numbers and deserves status alongside fertility, mortality, and immigration. The size of the aged population in 2050 is determined with arithmetic. Its functional health is not. If longevity gains arrive as compression of morbidity, disability concentrated into a shorter period before death, the 85+ boom produces far less care demand per person than headcount implies. If they arrive as expansion of morbidity, demand rises faster than the population does. The evidence is mixed: disability rates at given ages have improved, while dementia, obesity, and multiple chronic conditions push the other way.

The spread between the two paths is wide enough to move Medicaid long-term care projections, nursing home demand, home care hours, caregiver burden, and retrofit demand at once. This paper treats population aging as near-certain and healthspan as projection-dependent. The care-shortage argument holds under either path, because even compression leaves a far larger absolute number of people in the high-need years.

Two features of this uncertainty deserve emphasis, because it behaves differently from the other three variables in this section. First, no one forecasts it and publishes the number. Fertility, mortality, and immigration each carry an explicit assumption in the Trustees Report and the CBO outlook, stated and revised annually. Healthspan enters those projections implicitly, through per-beneficiary cost growth and disability incidence, where it is hard to isolate and easy to leave unexamined. A reader can check whether a fertility assumption is reasonable. Checking a morbidity assumption means reconstructing it.

Second, it is the one variable in this section that medicine could plausibly move. Fertility responds weakly to policy, mortality improves slowly, and immigration is a political decision. But a therapy that delayed the onset of dementia by five years, to take the clearest case, would do more to relieve the care constraint of Section 3.5 than any fiscal measure in this paper.

The possibility is real and it is not a reason for complacency, because the reverse is equally possible. Rising obesity and multimorbidity among the cohorts now entering their sixties could push the other way, and those cohorts are already formed. This is the largest genuine uncertainty in the care demand baseline, the one most likely to make the projections here look wrong in either direction. Immigration governs how many workers there will be; healthspan governs how much care the old will need, and the two are uncertain independently, and the one that receives the least attention in fiscal debate.

Aging works on the labor force through two channels, and only one of them is population size. The other is participation. Even if age-specific participation rates hold constant, the shift of population weight toward older ages drags aggregate participation down mechanically. Prime-age participation near multi-decade highs cannot offset the compositional pull.

Aging also transforms the dependency structure. The Social Security Trustees project the ratio of covered workers to Social Security beneficiaries, about 2.7 in the mid-2020s, falling toward 2.3 by 2040 and continuing down thereafter (Social Security Trustees 2026). **The same underlying shift, fewer workers and taxpayers relative to dependents, presses on systems as different as payroll-financed retirement benefits, enrollment-based school finance, and state and local legacy obligations, though the financing mechanics differ substantially in each case.**

Participation rates are behavior rather than arithmetic, and treating them as fixed would understate the economy's capacity to adjust. Employers facing labor scarcity have available responses that the projections do not assume: retaining workers past customary retirement ages, offering phased or part-time retirement, redesigning physically

demanding jobs, accommodating disability, and using experienced workers as trainers in the apprenticeship function Section 4.4 identifies as being at risk.

Older-worker participation has already risen substantially over three decades and could rise further. Two limits bound the offset. The response is available mainly to workers whose jobs can be held into the late sixties, which is the upper branch of the K-shaped economy, and it does not touch the care shortage, since the work requiring physical presence is the work older bodies are least able to perform. Higher participation among older workers is a real adjustment mechanism, but only a partial one.

2.2 Why Birth Rates Aren't Coming Back

Of everything in this section, the fertility decline is the hardest to reverse and gets the least attention. The aging story crowds it out, which is unfortunate, because the Social Security Chief Actuary concluded years ago that falling birth rates, not longer lifespans, are the dominant driver of rising program cost.

The total fertility rate fell from 2.12 in 2007 to approximately 1.6 by the mid-2010s and has not recovered. The 2025 provisional data extend the pattern: births fell one percent to 3,606,400, the general fertility rate fell to 53.1 births per 1,000 women aged 15-44, and the teen birth rate fell seven percent to another record low (NCHS 2026). CBO reduced its fertility assumptions in both its 2025 update and its 2026 outlook and now carries roughly 1.6 through 2055 (CBO 2026). The 2026 Trustees Report made a parallel downward revision.

Three features of the decline matter for everything that follows. First, it is concentrated among the young. Birth rates for women under 25 have fallen to record lows and keep falling, while rates for women over 30 have risen; the peak childbearing age group is now 30-34. This is delay, but delay at these magnitudes becomes forgone births.

Cohorts postponing from their twenties recover only a fraction in their thirties, and completed cohort fertility is drifting down accordingly. The distinction is visible abroad right now. South Korea’s fertility rate rose for a second consecutive year in 2025, to 0.80 from 0.75, on the largest increase in births since 2007; demographers attribute the rebound to a comparatively large cohort of women now in their early thirties and to marriages deferred during the pandemic, and note that the cohorts behind them are smaller (Statistics Korea 2026). A rise in the rate driven by favorable cohort structure and recovered timing is not the same thing as a rise in completed fertility, and this paper treats the two separately throughout.

Second, the decline has a K-shaped signature, and housing sits at the center of it. The steepest drops have occurred among the populations that define the lower branch: young workers facing housing costs that have outrun wages, childcare costs that exceed public college tuition in most states, and student debt service that competes directly with family formation. The chain runs from housing cost to delayed household formation to delayed or forgone births, which is why housing recurs throughout this paper as a demographic variable and not merely a financial one, particularly in Sections 5 and 8.

Fertility intentions surveys consistently find that Americans want roughly two children. Economic constraints appear to explain an important part of the gap between intended and achieved fertility, particularly among lower-income young adults, alongside changes in partnership, marriage, and preferences.

The cost of having a child in a lower-branch county now includes distance. More than 500 hospitals discontinued obstetric services between 2010 and 2022, 238 rural and 299 urban, against only 138 that added them, and by 2022 more than half of rural hospitals, 52.4 percent, no longer offered obstetric care, up from 43.1 percent in 2010 (Penn LDI 2025). Roughly 35 percent of American counties now qualify as maternity care deserts.

The mechanism is the fixed-cost problem that recurs throughout this paper: obstetric care requires around-the-clock staffing of physicians, nurses, and anesthesia regardless of birth volume, so a unit with falling deliveries cannot scale its cost down and closes. Falling births close the unit, and the closed unit raises the cost and risk of the next birth. This is Loop B of Figure 1 below operating on fertility directly rather than through the tax base, and it is the same arithmetic that Section 5.1 describes in school districts.

Third, the decline compounds. The small cohorts born after 2007 become the small cohorts of potential mothers of the 2030s, so births fall further even if the rate per woman stabilizes. This is demographic momentum operating in reverse. By the 2040s, smaller cohorts of potential mothers make reversal increasingly difficult even if fertility per woman rises. Policy can still shift the distribution of births geographically and by income, but the national total becomes progressively harder to move.

Can policy turn this around? The record offers little ground for optimism. Pro-natal cash incentives, like those used in Hungary, South Korea, and Japan, have primarily influenced when people have children rather than increasing the total number of births. Meanwhile,

more effective interventions, such as improving housing affordability, reducing childcare costs, and providing income stability for young adults, are the same policies that the K-shaped economy has consistently failed to deliver over the past fifteen years.

The baseline therefore treats 1.6 as a plateau rather than a trough, and treats any sustained recovery above 1.8 as an upside scenario requiring policy change of a kind not currently on the horizon. The immediate institutional consequence, the emptying of the school pipeline, is the subject of Section 5.

2.3 When Deaths Start Outnumbering Births

Deaths, meanwhile, rise on a schedule set by the age structure. Annual deaths, roughly 3.1 to 3.3 million in the mid-2020s, climb steadily as the boomer generation ages into the high-mortality years. The crossover point at which deaths exceed births nationally, an event that has never occurred in peacetime American history, has moved closer with each projection vintage. CBO's January 2025 outlook placed it at 2033. Its January 2026 outlook, incorporating the fertility and immigration revisions, places it at 2030 (CBO 2026). Once natural decrease arrives nationally, immigration is no longer one contributor to population growth among several. **It is the only mechanism for growth that remains.**

National mortality, however, is essentially an average of two different trends, and using it as a single national figure hides a crucial demographic fact. Chetty et al. (2016) matched 1.4 billion tax records to mortality data. Men in the top one percent of the income distribution at age 40 could expect to live to 88.3 years, longer than the average in any country on earth. Men in the bottom one percent could expect 72.7, a gap of 14.6 years for men and 10.1 for women. Between 2001 and 2014 the gap widened: the richest gained roughly three years and the poorest gained almost none.

At the county level, the divergence is comparable, with life expectancy a range of nearly twenty years across American counties. The Scandinavian rich-poor gap is four to ten years, depending on country (Kinge et al. 2019). The American gap is over fourteen.

This matters for the baseline in three ways a national mortality table cannot express. First, lower-branch counties enter natural decrease earlier than their age structure alone predicts, because higher age-specific mortality compounds out-migration and lower fertility; three mechanisms push in the same direction in the same places. Second, the deaths-of-despair pattern documented by Case and Deaton removes people from the prime-age labor force rather than from the retired population, making it a labor supply loss as well as a mortality event, concentrated in the counties Section 2.5 describes.

Third, any projection applying a national mortality improvement rate uniformly overstates longevity in lower-branch places and understates it at the top, which propagates into everything downstream from trust fund solvency to long-term care demand. Section 3.2 shows what this does to the incidence of a retirement age increase. My companion work develops the clinical dimension, including divergence in cancer stage at diagnosis and maternal mortality (Lane 2026c).

Natural decrease is not a future event at the local level. It is the present condition of a growing share of American counties, disproportionately the non-metropolitan, older, lower-branch counties of the K. For these places the national crossover date of 2030 arrived years or decades ago; **what the 2030s bring them is the loss of the offsetting in-migration that once cushioned it.** The county-level geography of natural decrease maps closely onto the geography of economic divergence documented in Lane (2026a), though the match is not perfect. The directional claim, that lower-branch counties face earlier and more severe aging than upper-branch metros, is robust to projection error in any single series.

2.4 Immigration: The Only Lever That Moves Fast

Everything above is arithmetic. Immigration is politics, which is why this paper carries scenarios rather than a projection. The recent record is a study in whiplash. Net immigration surged between 2021 and mid-2024 to two or three million a year at the peak, driven largely by humanitarian parole. It then collapsed. Suspended refugee admissions, terminated parole pathways, reduced visa issuance, and intensified interior enforcement produced a sharp contraction in 2025.

How sharp is genuinely unresolved, and the disagreement is instructive rather than a nuisance. Edelberg, Veuger, and Watson estimate that net migration fell to between negative 295,000 and negative 10,000, which would be the first negative year in at least half a century (Edelberg, Veuger, and Watson 2026).

CBO, working from a different method, estimates that total net immigration remained positive at roughly 410,000 in 2025, and describes that figure as a trough from which immigration rises steadily through 2030 (CBO 2026). The composition matters more than the total. Within that positive figure, CBO estimates that the other-foreign-national category, which covers parole entrants and people present without authorization, ran to negative 360,000, with 360,000 more leaving than arriving. The positive headline comes from lawful permanent residents and temporary visa holders offsetting a large negative elsewhere. The two estimates are therefore not simply methodological rivals. They weight

different populations, and Brookings is closer to the category where the policy actually bit. The Census Bureau's Vintage 2025 estimates, released in January 2026 with revised international migration methodology, report a historic decline without settling the question (U.S. Census Bureau 2026).

The estimates agree on direction and disagree on sign. The federal government does not observe this quantity directly, which explains part of the disagreement, and it is one reason this paper treats immigration as policy-contingent rather than as data.

My baseline follows CBO: a trough in 2025 and 2026, then recovery toward roughly one million net entrants per year and above. That baseline embeds a political forecast, that the current restriction regime is temporary, which this paper does not evaluate. Table 2 brackets the arithmetic under three plausible policy regimes rather than predicting which prevails. The differences between them are not marginal. They are the difference between a country of roughly 375 million in 2050 and one that has already begun to shrink.

They also carry a property no other variable here shares: **immigration policy is the only demographic lever capable of materially changing national outcomes within a single decade.** Fertility works on a twenty-year lag and mortality on a schedule. Migration works now, altering labor force growth and dependency ratios within a single business cycle.

Table 2. Immigration sensitivity scenarios, 2026-2050

Scenario	Net immigration assumption	Approx. 2050 population	Labor force and principal implications
Baseline (CBO)	Trough in 2025 (CBO: about 410,000; Brookings: negative), recovering after 2027 toward 1.0-1.2 million per year	~362 million	Labor force: slow but positive growth. Deaths exceed births from 2030; immigration accounts for all population growth thereafter

Low	Restriction persists; net flows near zero or negative through the early 2030s, then roughly 500,000 per year	~350-355 million, peaking in the 2040s	Labor force: flat, then contracting. Population decline begins before 2050; trust fund depletion accelerates; schools and housing demand fall fastest in lower-branch counties
High	Policy reversal restores flows near the 2010s norm of roughly 1.5 million or more per year by the 2030s	~375+ million	Labor force: clearly positive growth. National natural decrease is masked; dependency ratios improve modestly; fiscal gaps narrow but do not close

Sources: baseline from CBO (2026); alternative paths are author-constructed scenarios informed by the Census Bureau's low- and high-immigration projection variants, not direct reproductions of Census projections. Population figures are approximate midpoints. All three scenarios are policy-contingent by construction.

Table 2 states aggregate flows, and that framing conceals as much as it reveals. One million entrants is not demographically equivalent to another million. Age at arrival determines how many working years precede benefit claims and is the largest single driver of the fiscal effect. Skill mix determines which shortages are relieved. Whether dependents accompany workers determines school enrollment and the ratio of contributors to claimants. Settlement location determines which states gain representation.

A restrictive regime admitting younger workers with paths to permanence produces different arithmetic from a permissive one admitting older arrivals into insecure status. The question is not only how many, but who, at what ages, with what skills, and where, which is what makes the sector-specific channels of Section 9.3 analytically distinct from aggregate levels.

Two features of the sensitivity analysis matter. First, the scenarios are asymmetric in their labor market incidence. Restriction does not remove workers uniformly. It removes them first from agriculture, construction, food processing, hospitality, and direct care. Reduced labor supply in those sectors raises wages or cuts output, or both; costs rise for food, housing construction, and care services; and those cost increases feed back into the

housing and childcare constraints that suppress fertility, while landing hardest on lower-branch household budgets. A restriction regime therefore tightens the same care-economy labor markets that the 85+ boom is about to stress, a collision developed in Sections 3.5 and 4.

Second, the scenarios compound with fertility. In the high scenario, immigrant families partially offset native fertility decline, since immigrant fertility remains modestly higher and immigrants are concentrated in childbearing ages. In the low scenario, both engines fail at once, and the population pyramid inverts a decade faster than the baseline suggests.

It is worth asking, before leaving this section, what the low path would do across the systems this paper examines. The answer is not what a reader might assume, because the effects do not all run in the same direction, and the largest of them is not fiscal.

Three systems get unambiguously worse, and two of the three in ways that cut against the usual argument for restriction. Social Security's seventy-five-year deficit rises from 4.42 percent of taxable payroll toward roughly 5, by the Trustees' own sensitivity, as Section 3.2 works through. Medicare's Hospital Insurance fund deteriorates for the same reason and more sharply per capita: immigrants supplied about 14.7 percent of trust fund contributions against 7.9 percent of expenditures in the period Zallman and colleagues examined, a net subsidy from a population that pays payroll taxes for decades and claims comparatively little (Zallman et al. 2013). Restriction is frequently defended as reducing a fiscal burden. For both trust funds it does the opposite.

The third is the labor force, and it is the largest effect in the paper. Table 5 puts it at roughly 166 million in 2050 on the low path against 180 million on the CBO path, a gap of about 14 million workers. Everything in Sections 4 and 9 that depends on labor supply moves with that number.

Two systems accelerate rather than worsen. K-12 enrollment falls faster, since immigration supplies both children directly and a modestly higher birth rate. Higher education takes it twice, through domestic cohorts and through the international enrollment that Section 5.3 identifies as the sector's principal hedge. In both cases the low path does not create a new problem. It advances the schedule of one already underway.

Two are genuinely ambiguous, and saying so is more useful than guessing. Housing demand falls with fewer households forming, which relieves pressure in constrained metros, but construction labor is heavily immigrant and restriction raises the cost of building. The two effects partly offset and the net is unclear. Medicaid is ambiguous in dollars for a similar reason: fewer immigrant enrollees on one side, and on the other a

direct care workforce that is roughly a quarter foreign born, whose contraction raises the cost of delivering the same care.

That last case deserves separating out, because it is where the low path does the most damage and where the damage is least visible in a budget. The binding constraint identified in Section 3.5 is not money but people willing to do the work. Restriction reduces the supply of those people directly, at the same moment the 85-plus population begins its steep climb. Whatever it saves in enrollment it costs in capacity, and capacity is the thing that cannot be appropriated. Of everything the low path does, this is the effect that no fiscal projection will capture and no legislature will be forced to vote on.

2.5 Where People Are Going, and What They Leave Behind

Meanwhile, national near-stagnation coexists with vigorous internal churn, and the churn has a consistent direction. Growth concentrates in the South and Mountain West, in metropolitan areas generally, and within them in a subset of high-amenity, high-skill metros. Decline concentrates in the rural Northeast and Midwest, in small manufacturing-legacy cities, and in farm counties. In 2025, even with national net international migration near zero, states such as Florida and Texas continued to absorb both domestic migrants and the largest shares of what international migration remained (U.S. Census Bureau 2026).

One qualification belongs with every projection here that extrapolates the migration of the 2010s and early 2020s. Sun Belt growth is superimposed on rising exposure to heat, water stress, wildfire, and coastal flooding, and that exposure is arriving financially through property insurance, where premiums have risen sharply and carriers have withdrawn from some markets.

Affordability, insurability, and water availability are plausible counterweights, and Florida's already-slowing in-migration may be an early instance. This paper does not model climate migration. It notes only that the projections in Section 6.1 rest on an extrapolation, and this is one identifiable reason the extrapolation could fail.

The sorting is age- and skill-selective, which is what gives it its K-shaped character. Young, educated workers leave shrinking counties for growing ones, taking their future children and their peak taxpaying decades with them. The counties they leave age faster than either fertility or longevity alone would dictate. Out-migration removes the young directly and removes the next generation prospectively. The mechanism then feeds on itself: departures shrink the local tax base, a smaller tax base forces schools to consolidate and

hospitals to close service lines, and the withdrawal of those services pushes the next cohort of young families out, reinforcing the lower-branch profile.

By the 2030s the result is a demographic bifurcation layered on the economic one. Upper-branch metros experience national aging in diluted form, with in-migration continuously refreshing the working-age base. Lower-branch counties experience it in concentrated form: median ages a decade or more above the national figure, and a rising ratio of Social Security and Medicare income to earned income in the county economic base. In hundreds of counties, federal old-age transfers are already the largest single income source (Fikri, Eckhardt, and Glasner 2024). **This is the fiscal geography on which the trust fund politics of the early 2030s will play out**, and Sections 3 and 6 treat it at length.

2.6 Who the Country Is Becoming

Two more shifts complete the baseline. The first is generational. Millennials and Generation Z together already outnumber boomers and Generation X in the adult population, and by the mid-2030s they constitute a clear majority of the electorate and essentially the entire prime-age workforce. Their economic formation differs sharply from that of their predecessors. Entering adulthood into the 2008 crisis, the pandemic, and the housing cost escalation of the 2020s, they hold lower age-adjusted homeownership, form families later, and depend more on wage income relative to assets than earlier cohorts did at the same ages.

Calling these cohorts the lower branch of the K would be too simple. Older millennials are in their forties, many own homes with substantial equity, and within-generation dispersion is wide and widening. The accurate claim is narrower and more useful: their starting conditions have been less favorable at comparable ages than those of earlier cohorts, and the K runs through these generations as much as between them. A 34-year-old software engineer with vested equity in Austin and a 34-year-old warehouse supervisor renting in Youngstown share a generation and occupy opposite branches, and their household formation and fertility diverge accordingly.

What holds at the generational level is that the median member enters the 2030s with less cushion than the median boomer had at the same age, arriving at electoral dominance just as boomers reach peak entitlement dependence.

The second shift is racial and ethnic. Under Census Bureau projections, no single racial or ethnic group constitutes a majority of the population by the mid-2040s, and the crossover is already past for the under-18 population. The shift is generationally stratified: the older

population remains predominantly white non-Hispanic deep into the projection period while younger cohorts are already majority-minority. **The overlap of the age divide with the racial divide, an older, whiter population claiming age-based transfers financed by a younger, more diverse workforce, will shape coalition politics for two decades.**

And because both shifts are geographically clustered, they change not only the composition of the electorate but the distribution of political power across states. Section 6 develops the interaction with immigration policy, which slows the compositional shift only modestly while shrinking the financing workforce substantially.

2.7 These Problems Feed Each Other

The preceding subsections describe forces usually presented in sequence, as though demography caused economic outcomes which in turn caused fiscal outcomes. The actual structure is circular. Four reinforcing loops recur throughout this paper, and Figure 1 sets them out.

Loop A runs from sub-replacement fertility through a smaller entering workforce to fiscal pressure on age-based transfers, which may be resolved through higher taxes or thinner services. Whether that burden falls on young households, and whether it further suppresses fertility, depends on how the adjustment is designed. I treat Loop A as a plausible reinforcing mechanism rather than a demonstrated one. Loop B runs from the departure of young workers through rising median age and natural decrease to an eroding tax base, school and hospital contraction, and the exit of the families who remain.

Loop C runs from care wages held down by Medicaid reimbursement through worker shortage to higher costs and narrower access, shifting care onto unpaid family members whose withdrawal from the labor force cuts payroll revenue and raises eventual public liability. The return path is not automatic: falling federal revenue does not mechanically reduce state reimbursement rates. It closes only through budget decisions, since states facing weaker revenue and rising long-term care obligations have historically constrained provider payment rather than raised it. Absent that step the chain is a sequence rather than a cycle.

Immigration restriction is not part of that cycle; it is an external force that accelerates it. Loop D is regional rather than national. Fewer children in a region means falling enrollment and institutional revenue, campus closures, and shrinking local capacity to train nurses, aides, and technicians, which thins local services and pushes out the young families whose children would have filled the classrooms.

Why does this matter? Reinforcing loops behave differently from linear chains under policy. A linear problem can be addressed at many points along the chain with roughly similar effect. A loop can be entered at only a few points with real leverage, and interventions elsewhere are absorbed by the cycle and dissipate. The loops also explain why the K-shaped divergence widens rather than mean-reverting: each loop runs faster where absorptive capacity is thinner, so the same national trend produces accelerating separation. Sections 3 through 7 return to these loops by letter, and Section 9 asks which entry points actually carry leverage rather than merely relocating costs.

Figure 1. Four reinforcing loops in the American demographic system

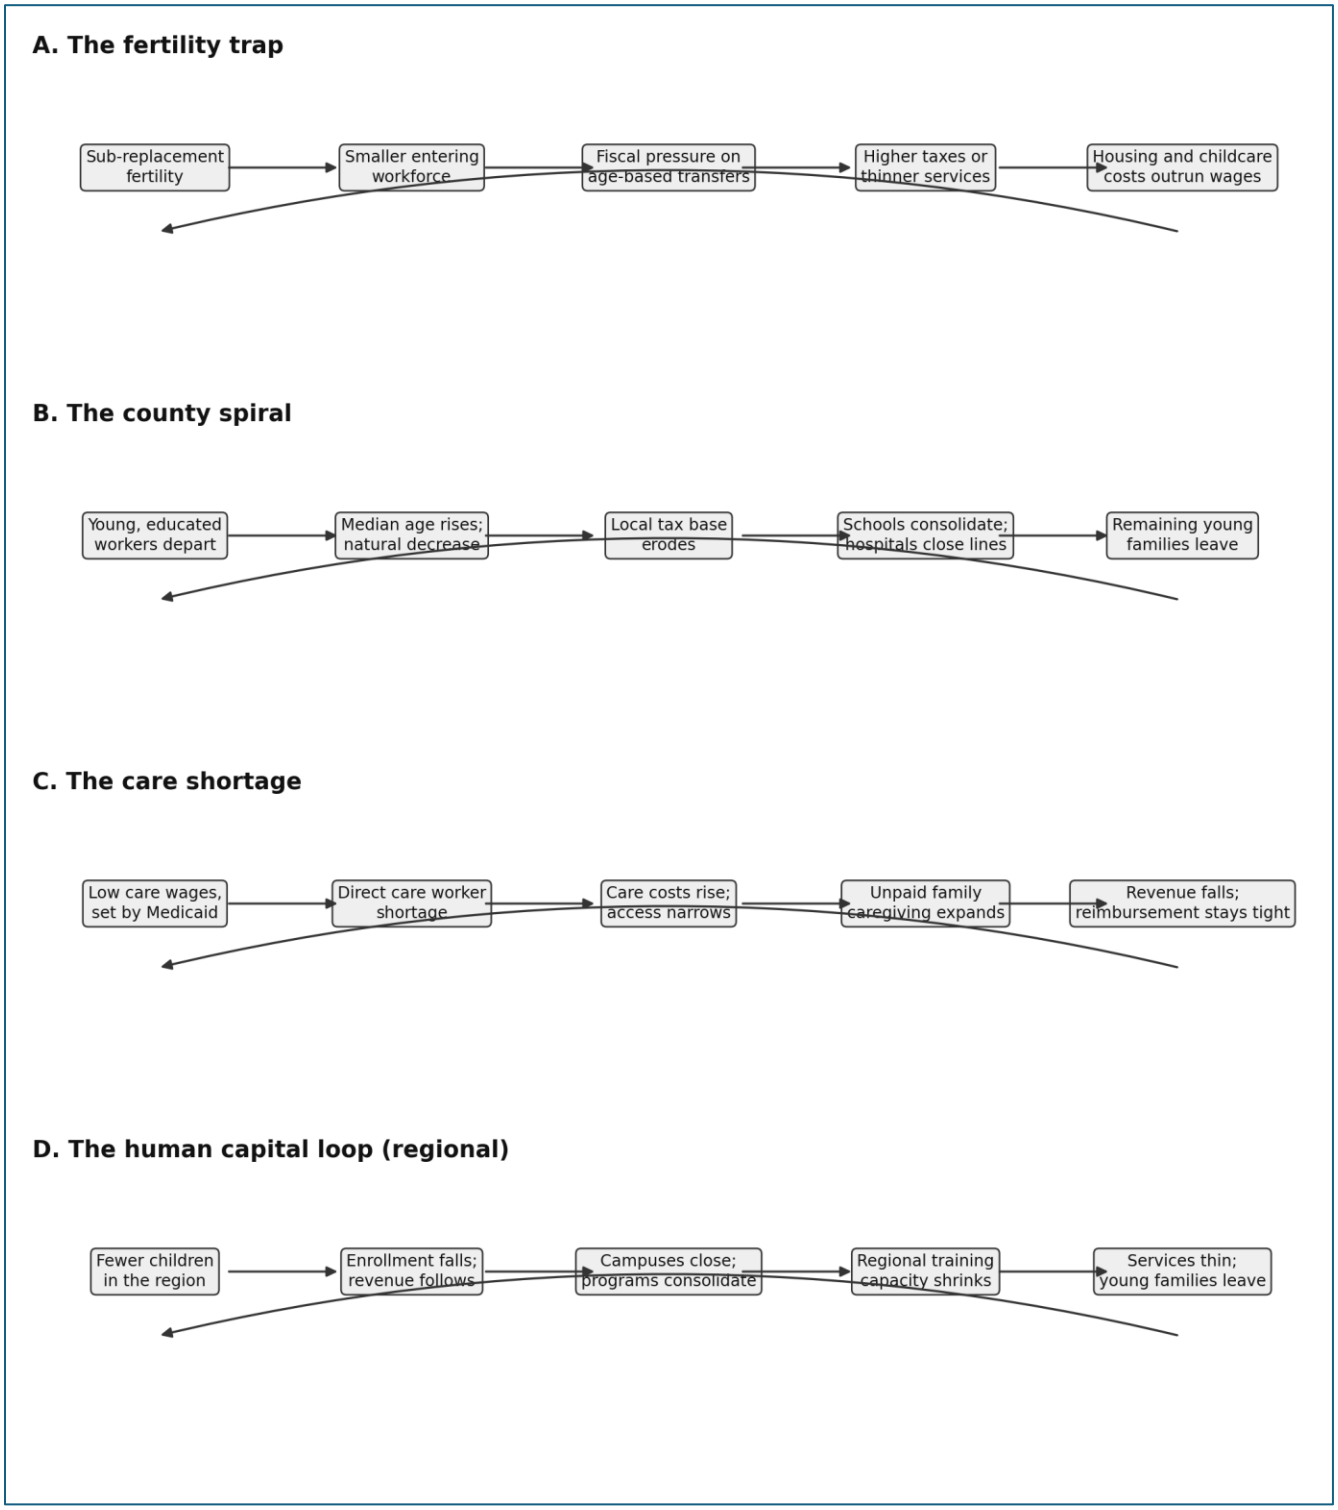

Source: author. Each arrow runs to the condition it worsens, and the return arrow closes the cycle back on its own starting condition. Loops A, B, and C operate nationally; Loop D operates at the regional level, since regional decline does not reduce national birth cohorts. Immigration restriction accelerates Loop C from

outside rather than forming part of it. Loop A is developed in Section 2.2, Loop B in Section 2.5, Loop C in Sections 3.4 and 3.5, and Loop D in Section 5.

2.8 Where My Framework Doesn't Work

A framework I apply to fertility, immigration, three entitlement programs, artificial intelligence, schools, colleges, reapportionment, and voting behavior is a framework at risk of explaining everything and therefore predicting nothing. So let me say plainly where it does not work.

Clearly, it does not predict where demographic shocks originate. Fertility decline is steepest in some of the most affluent metropolitan areas in the country, where housing costs have risen fastest. Some of the sharpest K-12 enrollment losses are in high-cost coastal cities that have priced out families with children, not in Appalachia. AI exposure concentrates among the credentialed, that is, among upper-branch workers. A framework that predicted lower-branch places would experience more of everything would be wrong on each of these.

What the framework does predict, and what the evidence in this paper supports, is absorptive capacity. San Francisco and a shrinking county in western Kansas can experience the same proportional fertility decline. San Francisco has universities that can be repurposed, hospital systems with diversified payer mixes, a tax base resting on assets rather than payroll, capital markets within reach, a state fiscal environment able to cross-subsidize local shocks, and continuing in-migration to replace departing residents. The Kansas county has one school district, one hospital operating on Medicaid margins, a tax base tied to a shrinking working-age population, and no in-migration. The shock is the same size. One place absorbs it and the other does not.

The K-shaped framework is a claim about the distribution of institutional and fiscal buffers, not a claim that lower-branch places receive larger shocks.

Two further limits should be noted. The framework is a two-branch simplification of a continuous distribution, and the middle is genuinely ambiguous; mid-sized metros with one large employer or one university do not sort cleanly. And the branches are not fixed. Places move between them, as several Sun Belt metros have over three decades. Where the paper uses the framework, it is making a claim about capacity to absorb, and where a phenomenon does not fit, this paper says so rather than stretching the category.

2.9 If the Door Stays Closed

Section 2.4 treats restriction as a path among three. This subsection asks a different question, and it is the one a reader is entitled to ask of a paper that makes immigration the swing variable. Suppose the current posture is not a two-year episode or a one-administration policy but the settled American position for the next quarter century. What kind of country does that produce?

Start with the arithmetic, which is the least contestable part. Under the low path the population peaks in the 2040s and begins to decline, which has not happened in American history. But the headline number understates it, because restriction compounds through a second channel. Immigrants arrive concentrated in childbearing years and have modestly higher fertility, so excluding them excludes their children as well. Section 2.2 shows how a fertility shortfall becomes structural once the cohorts that would have borne the additional births are themselves small. Sustained restriction adds to that shortfall every year it continues, and the compounding runs for as long as the policy does.

The geographic result is the sharper one, and it follows directly from Section 2.5. Immigration is not distributed evenly. Several large metropolitan areas depend on it for all of their population growth, and a number of counties that currently register as stable are stable only because international arrivals offset domestic departures. Remove the arrivals and those places join the category the paper describes as the lower branch. Indefinite restriction does not slow the K-shaped divergence. It widens it, by moving places from one side to the other.

The care constraint changes character rather than degree, and this is the consequence I would put first. If restriction is temporary, the shortage documented in Sections 3.5 and 4.2 is a difficult decade that ends when the pipeline reopens. If it is permanent, the shortage becomes a standing feature of American life, and the country must choose among three responses that it has so far avoided choosing among: pay substantially more for care, shift substantially more of it onto unpaid family members, or ration it. Those are the only options. Sustained restriction does not remove the need for several million additional caregivers. It removes one of the ways of meeting it.

There is a comparison case, and the paper should name it rather than imply it. Japan and Korea are what sustained low fertility combined with low immigration produces, and Section 9.7 examines their record on the fertility side. The demographic resemblance would be real. The institutional resemblance would not, because both countries pair that demography with universal health coverage, long-term care insurance, and household savings rates the United States does not have. A country choosing the same demographic path with a thinner set of public institutions is not running the same experiment.

Two things cut the other way, and leaving them out would make this section propaganda rather than analysis. Immigration at any plausible level cannot fully offset population aging. Pritchett estimates that holding working-age population ratios constant in aging rich countries through migration alone would require foreign-born shares far beyond anything a democracy has attempted, on the order of forty to fifty percent under his assumptions about accompanying dependents (Pritchett 2024). The magnitude depends heavily on which countries and which denominator are used, and I report it as an order of magnitude rather than a precise result. Immigration supplies part of the answer and cannot offset aging by itself, which is how the Migration Policy Institute puts it (MPI 2026). Restriction forecloses that half. It does not create the other half's problems, which fertility created on its own.

And the option may not stay open. The United States has treated immigration as a lever it can pull whenever it chooses, and that assumption is weakening. Fertility is falling almost everywhere: IHME projects that 155 of 204 countries will be below replacement by 2050, rising to 198 of 204 by 2100 (Vollset et al. 2024). Fertility among the traditional source populations for American immigration, including Mexico, has already collapsed. As working-age populations shrink across the rich world, countries accustomed to selecting among applicants will find themselves competing for a smaller pool. A nation that restricts today on the assumption that it can reopen later is making a bet not only about its own politics but about whether the people it declines to admit will still be available, and still interested, in twenty years.

That is the net effect. Not a single number, but a set of foreclosures: a population that peaks and turns down, a divergence that widens by reclassifying places, a care shortage that becomes structural rather than cyclical, and an option that closes quietly while the country assumes it remains open.

3. Three Programs, Three Different Failures

Three programs take most of what demographic change does to the federal budget, and they fail in three different ways. Social Security has a hard legal deadline in 2032 and a repair list that is relatively simple but politically radioactive. Medicare has a deadline too, but it covers the smaller half of the program while the larger half grows with nothing to stop it. Medicaid has no trust fund and therefore no deadline at all, which is exactly why the demographic wave hits it hardest and quietly. I take each in turn. Then I argue that the framing all three share, that these are problems of money, puts the constraint in the wrong place.

3.1 What a Trust Fund Date Actually Means

Start with what depletion is not. The OASI trust fund is an accounting record of past payroll tax collections in excess of past benefit payments, held as Treasury securities. Depletion means that record reaches zero. It does not mean the government runs out of dollars, cannot issue Treasury securities, or loses the operational ability to pay benefits.

What depletion means, precisely, is that a statute stops authorizing full payment. The Social Security Act permits the Administration to pay only from continuing income once reserves are exhausted, and, as a result, in the fourth quarter of 2032, revenue covers 78 percent of scheduled benefits (Social Security Trustees 2026). The cut is automatic because the law makes it so. And the law can be changed.

Section 1 laid out the two accounts of what follows. The conventional public finance view treats the shortfall as requiring higher taxes, lower benefits, or borrowing, each with real costs. The functional finance view holds that a currency issuer can always make the payments, so the choice among financing methods distributes claims on output rather than answering a question of affordability.

Both are right about something. Functional finance is right that the depletion date measures a legal authorization rather than a national capacity, and that prefunding financial claims does not by itself guarantee the real resources available when those claims are exercised, though financing choices can influence how much capacity is created and who controls it. The conventional view is right that financing methods are not interchangeable: general revenue, deficit expansion, and payroll taxation distribute purchasing power differently, and if claims on output outrun production, the adjustment arrives as inflation, itself a regressive tax.

Put this way, the frameworks differ less than the rhetoric around them suggests. Both end at inflation and real resource limits. A central disagreement concerns whether, and through what channels, financial prefunding today improves the real-resource position available tomorrow.

That relocation is the analytical value of the lens, and it survives disagreement about the rest of the framework. In 2040 the country will have some number of workers, some quantity of housing, and some number of people willing to bathe and feed the very old, and no arrangement of the ledger changes those quantities. A society that produces enough can support its elderly regardless of what the ledger says. A society that cannot produce enough will fail to support them no matter how carefully the ledger was maintained. **The dates in this section are real and will drive policy, because Congress wrote them into**

law and the political system responds to legal deadlines rather than to resource constraints. But they measure a deadline for political decision. Section 3.5 states what the underlying constraint actually is.

One statutory mechanism deserves separate attention, because it converts the accounting question of this section into a distributional one whether anyone intends that or not. The Statutory Pay-As-You-Go Act of 2010 requires the Office of Management and Budget to record the budgetary effect of every enacted law touching mandatory spending or revenue on two rolling scorecards, one covering five years and one covering ten. If either scorecard shows a net deficit increase when a session of Congress adjourns, the President is required to issue a sequestration order cutting non-exempt mandatory programs by enough to offset it (P.L. 111-139; CRS 2010). This is not a norm or a convention. It is law with an automatic enforcement mechanism, and it operates without a further vote.

What makes it consequential for this paper is the exemption list. Social Security is exempt from sequestration. So are Medicaid, SNAP, Supplemental Security Income, TANF, veterans' benefits, federal employee retirement, unemployment compensation, and refundable tax credits. Medicare is subject but capped at a four percent reduction (CRS 2010). The practical result is that roughly 235 of some 2,000 budget accounts are fully exposed, amounting to on the order of $150 billion to $200 billion out of a federal budget above $6.8 trillion, or something like two to three percent of the total.

Read carefully, that structure does not discipline existing spending at all. The programs this paper is about are largely outside the enforcement base. What the statute disciplines is new legislative action: any bill that increases the deficit adds to a scorecard whose consequences fall on a narrow set of accounts that had nothing to do with it.

And Congress routinely disables even that. The standard practice is to write a provision into the bill directing that its effects not be entered on the PAYGO scorecards, which is why the accumulated scorecard balance has run into the trillions, far exceeding what the sequestrable base could ever offset, without triggering the cuts the law appears to require. Where a bill cannot carry such a provision, as with a reconciliation measure, Congress has passed separate relief, as it did after the American Rescue Plan Act.

So the constraint binds selectively, and the selection is the point. A measure that reduces claims on output generates its own offset and needs no waiver. A measure that expands productive capacity, raising care wages or funding the training pipelines of Section 5.4, requires either an offset or a majority willing to suspend the rule. Under this architecture, the two are not symmetric, regardless of which produces more real output over twenty years. The mechanism is neutral in its stated form and conservative in its effect, in the

literal sense of favoring the existing distribution, and Section 2 established that the existing distribution is the K-shaped one.

The 2025 reconciliation act supplies a live illustration and a warning about how these mechanisms compound. Because the act increases the deficit and because Medicaid and Social Security are exempt, its scorecard effect falls on Medicare. Published projections put the potential statutory PAYGO reductions in the hundreds of billions of dollars across the 2026 to 2034 window under current law and absent congressional intervention, with the first year in the tens of billions. Congress has repeatedly neutralized such sequestrations, so these are contingent magnitudes rather than expected outcomes.

The specific magnitudes are contested and depend on scoring assumptions, so I report the order of magnitude rather than a point estimate. That would be a benefit reduction arriving not because anyone voted for it but as the automatic consequence of a tax and spending bill, and it would stack on top of the separate eleven percent Part A cut that Section 3.3 describes at Hospital Insurance depletion. Two independent statutory triggers, two automatic cuts, neither requiring a vote. Section 9.9 returns to this when it asks why the interventions with the most leverage have the least political room. I have argued the sharper version of this case elsewhere, in the specific context of Social Security financing (Lane 2024).

3.2 Social Security: Six Years to Decide

The 2026 Trustees Report moved OASI depletion to the fourth quarter of 2032, one quarter earlier than the prior year. The combined OASDI date held at 2034, but combining the funds requires an act of Congress, so the operative date for retirees is currently projected to be 2032, not 2034.

At depletion, continuing revenue covers 78 percent of scheduled OASI benefits, and the payable share drifts down toward 62 percent by 2100 (Social Security Trustees 2026). The 75-year actuarial deficit is 4.42 percent of taxable payroll, up from 3.82 percent the year before, and the 75-year unfunded obligation rose from $25.1 trillion in the 2025 report to $29.3 trillion in the 2026 report.

Those figures assume the immigration path of Section 1.1. It is worth asking what the low path would do to them, and the Trustees publish enough to answer. Their sensitivity analysis reports that raising average annual net immigration by 100,000 improves the long-range actuarial balance by roughly 0.10 percent of taxable payroll, and their intermediate assumption is average net immigration of about 1.2 million a year over the projection period. The low path of Section 2.4 runs roughly 700,000 below that. Applying the Trustees’

own elasticity gives about 0.7 percent of payroll, which would put the 75-year deficit near 5.1 percent rather than 4.42. It is a rough calculation and I present it as one, but it uses the Trustees' published sensitivity rather than an assumption of my own.

The more surprising result is what the low path does not do. The Trustees' sensitivity analysis finds the projected depletion date essentially unchanged across their low, intermediate, and high immigration assumptions. Immigration moves the seventy-five-year balance materially and the depletion date barely at all, because depletion in the early 2030s is driven by a retirement wave already in motion rather than by the size of the entering workforce. This matters for how the policy discussion in Section 9.3 should be read. Immigration is the fastest lever available for labor supply, and it is not a solution to 2032. Anyone offering it as one is answering a different question than the trust fund date poses.

That $4.2 trillion revision deserves careful reading, because it is easy to overstate. About $1 trillion of it comes from moving the valuation window forward a year rather than from any change in the underlying outlook. It is not the price tag of one year of restrictionist immigration policy. It is a change in a 75-year present-value estimate produced by revising several assumptions at once: lower fertility across the whole projection horizon, lower immigration across the whole projection horizon, the reduced taxation of benefits enacted in the 2025 reconciliation act, and economic and methodological updates.

What the revision does establish is the extraordinary sensitivity of long-range fiscal projections to demographic parameters. Modest downward adjustments to two demographic assumptions moved the program's long-run shortfall by an amount comparable to a major tax or benefit reform. **Demographic policy is fiscal policy, which is why a paper about population change is necessarily a paper about the federal budget.**

So what are the options? The list is short and well understood. Table 3 sets out the main ones with their approximate contribution to closing the 4.42 percent gap and an assessment of who bears the adjustment. The distributional column resists one-word labels: each option's incidence depends on interactions with the benefit formula, the taxable maximum (the wage ceiling above which earnings are not subject to payroll tax), and the financing method.

Table 3. Selected Social Security repair options and their distributional incidence

Option	Share of the 4.42 percent gap closed	Who pays	Distributional assessment
Eliminate or raise the taxable maximum	Roughly 60 to 70 percent, depending on whether the new earnings earn benefit credit	Earners above the cap, concentrated in upper-branch metros	Progressive on the revenue side. Draws from the branch whose income has grown fastest since 2020. Granting benefit credit on the new earnings returns part of the revenue as higher benefits to high earners
Raise the full retirement age to 69 or 70	Roughly 30 to 40 percent	All future retirees, as an across-the-board benefit cut	Regressive in effect rather than by formula. Lower-branch workers have shorter life expectancy at 40 and hold jobs that cannot be worked into the late sixties, so they collect for fewer years and cannot offset by delaying
Switch to chained CPI for the COLA	Roughly 15 to 20 percent	Current and future beneficiaries, with the cut compounding by year of receipt	Regressive in effect. The compounding cut is largest for the oldest beneficiaries, who hold the least non-Social-Security income and face the highest medical cost share

Raise the payroll tax rate	An immediate rise from 12.40 to 16.65 percent, 4.25 points, restores 75-year solvency	All wage earners from the first dollar	Mixed. The tax is regressive within wages and exempts capital income, but the progressive benefit formula returns more per dollar to low earners. Net incidence depends on whether the cap moves at the same time
Fund the gap from general revenue	Any share, by construction	Depends entirely on the financing method	Indeterminate ex ante. Incidence depends on the taxes, borrowing, spending offsets, or accommodation used. Removes the statutory trigger without resolving the real resource question

Gap-closing shares are approximate and drawn from the Social Security actuaries' provisions library; distributional assessments are the author's. Combinations are the realistic path, since no single option is likely to pass alone.

The retirement age option deserves scrutiny, because it is most often described as a technical adjustment to longevity and is in fact the option whose incidence diverges most sharply from its formal design. On paper it applies a uniform reduction. In effect it does not, because expected collection periods differ enormously by income, as Section 2.3 established. A lower-branch worker who spent thirty years on his feet and dies at 73 loses a far larger share of his expected lifetime benefit than a professional who works from a desk to 70 and lives to 87. He is also least able to respond by delaying, because the jobs that wear out bodies cannot be held into the late sixties.

Raising the retirement age is sold as a technical adjustment to longevity. It is a benefit cut aimed with unusual precision at the people who die earliest.

And the window closes on a schedule. Reforms of this program have historically required long phase-ins to avoid abrupt changes for people near retirement; the 1983 amendments phased the retirement age increase over four decades. Every year Congress waits shortens the available phase-in and narrows the menu toward options that take effect immediately, which in practice means options that hit current beneficiaries.

The usable window runs from roughly 2027, after the 2026 midterms, to 2032. Six years is enough time in principle; the 1983 fix took a commission and about two years from formation to enactment. The arithmetic of phase-in time is not a political prediction: the longer the delay, the more the feasible menu narrows toward options falling on current and near-term beneficiaries, independent of anyone's intentions.

So who actually absorbs a 22 percent cut? Not everyone equally. In hundreds of counties, mostly rural, mostly in the lower branch, Social Security and Medicare transfers are already the largest single source of personal income. Section 2.5 described how those counties formed. A benefit cut in such a place is not a household budgeting problem. It is a contraction in the county's aggregate income, transmitted immediately to local retailers, health providers, and property tax collections, because retirees spend nearly all of what they receive and spend it locally. Loop B of Figure 1 runs in reverse and faster. Section 6 takes up the political consequence.

3.3 Medicare: Everyone Is Watching the Wrong Number

Medicare's Hospital Insurance trust fund is forecast to be depleted in the second quarter of 2033, one quarter earlier than last projected, with revenue covering 89 percent of costs in 2033 and 87 percent in 2034 (Medicare Trustees 2026). That leaves an aggregate financing shortfall of 11 percent, widening toward 16. How the shortfall would be distributed across providers and claims is not specified by statute.

Part A is not a small program. It runs about 1.5 percent of GDP and rises toward 2.0 percent by 2050, and an 11 percent payment cut would land hard on hospitals, skilled nursing facilities, and home health agencies, particularly the rural providers already operating at the margin. The HI deadline is not unimportant. It receives disproportionate attention because it comes with a statutory trigger, while the larger long-run exposure sits in the parts of Medicare that have none.

The national context matters for interpreting these figures. Total U.S. health expenditure reached 18.0 percent of GDP in 2024 and is projected by CMS actuaries to reach roughly 20.3 percent by 2033, growing consistently faster than the economy. Medicare's trajectory is a component of that system-wide pattern rather than an independent phenomenon,

which is the first reason to be skeptical of demographic explanations for it: the same excess growth appears in the commercial market, where the population is not aging.

Parts B and D cannot have a depletion date, because they are financed by beneficiary premiums and general revenue set annually to cover the following year's expected costs. They are solvent by construction, which means the constraint on them is political rather than statutory. And they are where the growth is.

In 2025, Part B spending reached $584 billion, roughly $140 billion more than Part A. The Trustees project Part B growing 9.5 percent annually and Part D 9.4 percent annually through 2030, against roughly 4 percent nominal GDP growth (Medicare Trustees 2026). Part D was revised up sharply this year on GLP-1 and specialty drug utilization, with spending projected to nearly double from $181 billion in 2025 to $346 billion in 2035 (KFF 2026a). Total Medicare spending rises from 3.9 percent of GDP in 2025 to 6.5 percent by 2050. The HI trust fund conversation does not address that trajectory; it addresses a statutory trigger attached to a smaller component of the program.

Two things matter here. First, demography and pricing interact multiplicatively rather than additively. The aging of the boomer cohort sets the enrollment path and therefore the volume of services. Per-beneficiary cost growth, driven by service intensity, drug prices, technology adoption, and Medicare Advantage payment dynamics, sets the per-unit cost.

Demography explains enrollment growth; it does not explain the full spending path, and an analysis that treats Medicare cost growth as purely a population phenomenon will prescribe benefit cuts for what is substantially a per-beneficiary cost problem involving prices, utilization, technology, and payment design.

Second, the financing shift has a statutory tripwire already pulled. When general revenue is projected to exceed 45 percent of Medicare financing within seven years, the Trustees must issue a funding warning. The 2026 report projects that threshold is crossed in fiscal 2026, the tenth consecutive report containing such a determination and the ninth consecutive year in which two consecutive determinations have triggered the statutory funding warning (Medicare Trustees 2026). **Nine consecutive warnings producing no legislative response show what these triggers are. They are information, not constraint.**

The distributional incidence of Medicare's trajectory runs through premiums and cost sharing rather than through eligibility. Part B and Part D premiums consume roughly 12 percent of the average Social Security benefit today and are projected to reach nearly 21

percent by 2100, with total out-of-pocket costs approaching 41 percent of the average benefit (Medicare Trustees 2026).

For an upper-branch retiree with a 401(k), a paid-off house, and appreciated equities, a rising premium is an annoyance. For a lower-branch retiree whose Social Security check is 80 or 90 percent of total income, rising premiums and cost sharing convert Medicare's fiscal trajectory into an immediate consumption constraint. The program remains universal in name. The experience of it diverges by branch, and diverges further every year.

3.4 Medicaid: The Program That Gets Hit Hardest

Medicaid gets the least attention in these discussions and deserves the most. Three reasons. It is the nation's primary payer for long-term care, covering roughly 44 percent of institutional long-term care costs and the large majority of nursing home residents (Chidambaram et al. 2025). It is financed jointly by states that must balance their budgets, so it cannot run deficits when demand rises. And it faces the steepest demographic gradient of the three programs, because its cost driver is not the 65+ population but the 85+ population, which nearly triples from 6.5 million in 2022 to 17.5 million by 2060 (PHI 2025).

Importantly, the timing matters. Boomer retirement, which dominated the policy conversation of the 2010s and 2020s, is a Social Security and Medicare event. Boomer frailty is a Medicaid event, and it arrives later: the mid-2030s through the 2040s. **The country is currently designing Medicaid policy for a demographic moment that has not arrived, using the assumptions of a moment that is passing,** because the policy cycle runs on two-year and four-year horizons while the demographic cycle runs on decades.

And we are cutting it. The 2025 reconciliation act reduces federal Medicaid spending by an estimated $911 billion between 2025 and 2034, roughly 14 percent of federal funding for the program (KFF 2025). The mechanisms are work requirements for certain expansion adults beginning January 1, 2027, with several states implementing earlier, more frequent eligibility redeterminations, cost sharing up to $35 per service for certain expansion adults from October 2028, a provider tax freeze with the expansion-state cap falling from 6 percent to 3.5 percent by 2032, and limits on state-directed payments. CBO attributed the largest single share of projected coverage loss to the work requirements alone.

The mechanism reaching long-term care is indirect and easy to miss. Home and community-based services are optional Medicaid spending, and optional spending is most of what states can legally cut. Much of it already runs under capped waiver authorities with

waiting lists, so contraction bites first through lower caps and longer waits rather than visible terminations.

When federal funding falls and states cannot raise the provider taxes financing their share, the arithmetic pushes them toward exactly these categories. The precedent is unambiguous: during the last major reduction in federal Medicaid spending, 40 states served fewer people through home care and 47 cut benefits or payment rates (KFF 2025).

Cutting home and community-based services does not reduce the need for care. It shifts care to unpaid family members, and the scale of that transfer is often missing from fiscal discussions of aging. AARP estimates that 59 million Americans provided care to an adult in 2024, delivering 49.5 billion hours of unpaid care worth roughly $1.01 trillion at replacement wages, equivalent to nearly 24 million full-time workers, or about 17 percent of the full-time workforce (Houser et al. 2026). That exceeds total federal, state, and local Medicaid spending in the same year. **The unpaid caregiving sector is already larger than the paid one, and it is the shock absorber that every proposal to reduce public long-term care spending implicitly relies on.**

This is Loop C of Figure 1, and it closes on the fiscal system it was meant to relieve. Caregivers are disproportionately women in their fifties and sixties, at their peak earning years and within a decade or two of their own retirement. **Care obligations reduce their hours, their labor force participation, and their earnings, which reduces payroll tax revenue now and their own Social Security benefits later.** The federal savings are real at the point of appropriation. A substantial share is recovered as costs elsewhere: lower revenue, higher eventual benefit need, and the institutional placements that occur when home care becomes unavailable and cost more than the home care would have.

The geographic incidence completes the picture. Rural hospitals and nursing facilities operate on Medicaid margins in exactly the lower-branch counties described in Section 2.5, where the payer mix is worst and the population oldest. Roughly 136 rural hospitals have closed outright since 2010, and the pace has been faster in states that declined Medicaid expansion. KFF estimates the 2025 reconciliation act reduces federal Medicaid spending in rural areas by roughly $137 billion over ten years, against $50 billion appropriated over five years for the Rural Health Transformation Program. KFF cautions that the two figures are not directly offsetting, since most of the Medicaid reductions take effect after the rural fund is exhausted (KFF 2025).

A percentage cut applied uniformly at the federal level therefore lands as a solvency question in those places and as a budget line in wealthy suburban systems. The consequences are not confined to health care. A rural hospital is frequently the largest

employer in its county, so closure produces job loss, population decline, and falling property values alongside the loss of emergency care, and the counties losing hospitals are the same counties consolidating schools and losing colleges. Loop B again, and current Medicaid policy is accelerating it.

3.5 You Can't Appropriate Caregivers

Return to the framing argument with the numbers in hand. Suppose Congress solves every accounting problem in this section. It eliminates the taxable maximum, funds Part A from general revenue, and restores the Medicaid cuts. Every trust fund date disappears.

What has not changed is the economy's capacity to produce the care the 85+ population will require. Between 2024 and 2034, the direct care sector must fill an estimated 9.7 million job openings, more than any other occupational category, including roughly 772,000 net new jobs, with the remainder from turnover (PHI 2025). Median hourly pay was $17.36 in 2024 and median annual earnings just under $26,000, which is why turnover generates most of the openings. Roughly a quarter of direct care workers are foreign born, and at least one in five immigrant direct care workers comes from a country covered by the current visa pause (KFF 2026b).

Capacity and headcount are different things, and the difference matters for what follows. Capacity is the product of workers, the capital and technology they work with, the efficiency of the delivery system, and the degree to which one input can substitute for another. Several of those terms can rise. Care delivery could become more efficient. Assistive technology, remote monitoring, and eventually some forms of robotics could raise output per caregiver.

Section 4 takes up artificial intelligence at length, and a paper treating the worker-to-retiree ratio as destiny would contradict itself two sections later. Productivity growth can relieve part of this constraint. **Betting the retirement security of 80 million people on an unspecified productivity rate in a sector defined by physical presence is a poor plan.**

Because that is the distinctive feature of direct care: it resists the substitutions that have relieved labor constraints elsewhere. Someone must be in the room to lift a person out of bed. The tasks that dominate the work, bathing, feeding, transferring, and supervising, have proven durably resistant to automation, and the sector's measured productivity growth has been near zero for decades for exactly that reason.

This is Baumol's cost disease in its purest form (Baumol and Bowen 1966), and it means the care sector will absorb a rising share of national labor input as a matter of arithmetic

rather than policy failure. You can appropriate dollars by statute. You cannot appropriate caregivers.

Put the pieces together. The 85+ population climbs steeply from the mid-2030s. The workforce that cares for it is underpaid, has high turnover, draws heavily on immigrants, and is losing its immigrant pipeline under the low-immigration scenario of Section 2.4. The domestic replacement pool is shrinking because of the fertility decline of Section 2.2.

Medicaid, the payer that effectively sets wages in this sector, is being cut. And the unpaid family sector that absorbs the overflow is itself constrained by the smaller sibling networks that sub-replacement fertility produces, since the adult children available to care for each aging parent grow fewer with every cohort. Five trends converge on the same shortage.

This reframing changes what counts as a solution. Under the conventional frame, the policy question is how to reduce the dollar cost of old-age programs. Under the capacity frame, the question is **what would actually expand the supply of care: compensation high enough to retain workers, training and credentialing pipelines, immigration channels sized to the sector's demand, technology that raises output per caregiver where the work permits it, and possibly direct public employment.** Each has costs and failure modes, and Section 9 evaluates them against one another rather than assuming an answer here.

What Section 3 establishes is only that **the question has been misspecified**. The trust fund debate asks how to make the ledger balance. **The demographic evidence says the harder problem is filling several million projected openings with people willing to do difficult work for wages the current system will not pay.**

3.6 What Aging Does to Asset Prices

One consequence of population aging falls entirely outside the trust fund frame and belongs here because it bears directly on it. Older households hold more assets than younger ones, so an aging population mechanically raises aggregate demand for assets.

Auclert, Malmberg, Rognlie, and Straub (2025) formalize this as a race between asset supply and asset demand, and find that from 1950 to the present aging, rising inequality, foreign demand for U.S. assets, and slowing productivity growth together produced **a surge in asset demand large enough to outpace the growth in asset supply, which is a substantial part of why real interest rates fell** over that period.

Their projection to 2100 reverses the result, and the reversal matters for this paper in two directions. Aging continues to drive up asset demand. But at current tax and benefit levels, rising entitlement costs push government debt up faster, so asset supply wins. **They calculate that strong asset demand creates room for federal debt to reach roughly 250 percent of GDP without higher interest rates**, while stabilizing debt at any level would require a permanent fiscal adjustment of roughly 10 percent of GDP.

That second number is the most serious quantitative challenge to the framing adopted here. Section 3.1 argues that the constraint on old-age provision is productive capacity rather than financing. A permanent adjustment of 10 percent of GDP is a claim about magnitude, not about solvency, and it is compatible with the functional finance position: it says the real resources required to honor current commitments substantially exceed what current tax and benefit parameters allocate to them.

Read that way, the finding supports the argument of this section rather than undermining it. What it forecloses is the comfortable version, in which recognizing that a currency issuer cannot be forced into default is mistaken for a demonstration that the commitments are affordable at current settings. They are not, and the adjustment is large.

The distributional consequence closes the loop with the rest of the paper. If aging raises asset demand and supports asset valuations, the gain accrues to the households that own assets, which Section 2.6 and Section 8 identify as the upper branch. The same demographic force that strains public old-age provision, felt most acutely by households dependent on transfers, simultaneously supports the portfolios of households that are not. **Aging is not distributionally neutral even in its financial market effects, and the direction of the effect widens the K.**

3.7 Where That Leaves Us

The deadlines are legal triggers rather than measures of national capacity, but they will drive policy anyway, and the arithmetic of phase-in time means that delay narrows the feasible menu toward options falling on current and near-term beneficiaries. Second, demographic parameters now do more work in long-run fiscal projections than most conventional macroeconomic scenarios do, and the $4.2 trillion revision illustrates the magnitude.

Third, **the constraint that ultimately binds is productive capacity in the care sector, where measured productivity growth has been near zero for structural reasons**. Every current policy trend pushes that capacity the wrong way at once: immigration restriction,

Medicaid reimbursement cuts, and the fertility-driven contraction of the young adult population. Section 4 takes up the labor market directly.

Table 4. Fiscal dates established in this section

Date	Milestone
FY 2026	General revenue crosses 45 percent of Medicare financing; tenth consecutive excess-general-revenue determination, ninth consecutive funding warning
Dec. 2026	Medicaid work reporting requirements take effect
Oct. 2028	Medicaid cost sharing begins for certain expansion adults
2027-2032	The usable window for Social Security legislation with meaningful phase-in
2032	Expansion-state Medicaid provider tax cap completes its decline to 3.5 percent
Q4 2032	OASI reserves deplete; 78 percent of scheduled benefits payable
Q2 2033	Medicare HI reserves deplete; automatic cut of 11 percent rising toward 16
2034	Combined OASDI depletion if Congress merges the funds; 83 percent payable
Mid-2030s	The 85+ population accelerates and Medicaid long-term care demand climbs
2050	Medicare reaches 6.5 percent of GDP, up from 3.9 percent in 2025

Dates reflect current law and the intermediate assumptions of the 2026 Trustees Reports. Confidence classes and sources for each entry appear in the consolidated timeline, Table 12 in Section 10.

4. Two Shocks, Moving in Opposite Directions

American labor markets are taking two structural shocks at once, and the public conversation has collapsed them into a single story about technology and jobs. **They are not one story**. Artificial intelligence is reducing demand for certain tasks and appears to be reducing demand for some categories of labor. Demography is reducing the supply of labor overall and creating acute shortages in the kinds of work AI cannot do. These forces do not cancel out, because they act on different workers, in different occupations, on different timetables, and in different places. A country can run a serious shortage of home health aides and a serious surplus of junior analysts at once.

4.1 The Workforce Stops Growing

Start with the supply side. The Bureau of Labor Statistics projects labor force participation will fall from 62.6 percent in 2024 to 61.1 percent in 2034, a decline of 1.5 percentage points, representing roughly 4.3 million fewer participants than the 2024 rate would have produced (BLS 2025).

The composition of that decline matters more than its size. Prime-age participation remains near multi-decade highs and is projected to drift down only slightly. The largest single contributor is workers aged 16 to 24, whose numbers fall because the cohorts are smaller and their attachment declines accordingly. The rest is compositional: an aging population moves more people into age brackets with structurally lower participation, and no plausible increase in older-worker participation offsets the shift in weights.

So what does that do to growth? The arithmetic is standard. Output growth is the sum of growth in labor input and growth in output per unit of labor input, where labor input reflects hours and workforce quality, as well as headcount. As labor force growth approaches zero, which the low path reaches and the baseline approaches, maintaining historical rates of trend output growth requires substantially faster growth in productivity, hours worked, labor quality, or capital intensity than the postwar record delivered. None of those is impossible.

But a projection that assumes historically normal output growth while also accepting near-zero labor force growth has quietly assumed that productivity will accelerate enough to replace the missing labor. That assumption is rarely stated because it is not entered directly; it follows from two other assumptions made separately. And because higher assumed growth improves every fiscal measure, revenue, debt ratios, and depletion dates alike, a projection carrying it unstated is not merely uncertain**. It is biased toward the conclusion that the problem is manageable.**

International bodies have reached the same conclusion based on data outside the American data. The OECD projects that aging alone, holding labor productivity constant, will cut GDP per capita growth across the OECD area by about 40 percent, from 1.0 percent per year in 2006-19 to 0.6 percent per year over 2024-60. Over the same period, the working-age population falls by 8 percent, and the old-age dependency ratio rises from 31 percent to 52 percent (OECD 2025a).

One qualification applies to that figure and cuts against overreading it here: the OECD expects the United States to be among the countries in which GDP per capita growth declines the least, precisely because labor input had already contributed little to American growth. The American version of this problem is therefore less a coming deceleration than the continuation of one already underway.

The immigration scenarios in Section 2.4 bear directly and asymmetrically on this. Under the baseline, labor force growth is slow but positive. Under the low scenario, it flattens and then declines during the 2030s, marking the first sustained contraction of the American labor force in the country's history. Under the high scenario, it stays clearly positive. **No other policy lever available to Congress moves labor supply on this scale within a decade, which is worth stating plainly because immigration is almost never debated as labor market policy.**

Table 5 puts approximate numbers on that difference. It is a stylized calculation rather than a published projection, and the assumptions are stated so a reader can substitute their own. It takes the Bureau of Labor Statistics participation path as given, applies it to a domestic working-age population fixed by cohorts already born, and varies only the immigration contribution across the three scenarios of Section 2.4. Roughly three-quarters of immigrants arrive between ages 18 and 64, so the labor force effect of a given flow is larger and faster than its effect on total population.

Table 5. Labor force under three immigration scenarios, illustrative sensitivity exercise

	Low immigration	Baseline (CBO)	High immigration
Assumed net flow after 2027	about 500,000 a year	about 1.0 to 1.1 million a year	about 1.5 million a year
Labor force, 2030	roughly 170 million	roughly 171 million	roughly 172 million

Labor force, 2040	roughly 170 million	roughly 176 million	roughly 181 million
Labor force, 2050	roughly 166 million	roughly 180 million	roughly 191 million
Average annual growth, 2026-2050	about zero, turning negative in the 2040s	about 0.25 percent	about 0.5 percent
Difference from baseline in 2050	about 14 million fewer workers		about 11 million more workers

Author's stylized calculation. Method: 2024 labor force of about 168 million from BLS; participation declining from 62.6 percent in 2024 to 61.1 percent in 2034 and drifting down thereafter; domestic working-age population from cohorts already born through 2044, and from the fertility assumptions of Section 2.2 thereafter, which makes the 2050 column more projection-dependent than the earlier ones; immigrant arrivals assumed 75 percent of working age at entry, entering the labor force at the prevailing participation rate. Figures are rounded to the nearest million and should be read as magnitudes rather than forecasts. The purpose is to show the spread between policy regimes, not to predict a level.

The spread is the point. Under the low scenario, the American labor force stops growing and then contracts, which has never happened outside wartime. Under the high scenario, it continues to grow at roughly half the postwar rate. The gap between the two by 2050 is on the order of 25 million workers, which is roughly the size of the combined labor forces of Texas and Florida. No other policy lever available to Congress moves a quantity of that size within the projection period.

4.2 Where the Jobs Actually Are

Set the direct care shortage established in Section 3.5 against Section 4.1, and the problem states itself. **The occupation with the largest projected labor demand in the United States pays wages that keep its workers near poverty, draws roughly a quarter of its workforce from immigrants at a moment of restriction, and must recruit from an entering cohort that shrinks every year.**

The usual market response to shortage is rising wages, and wages in the sector have risen by about $2.98 an hour in real terms over the decade to 2024. The market cannot fully clear through wages, however, because its dominant purchasers constrain reimbursement. Roughly half of long-term care is paid for by Medicaid, whose rates are set administratively

by states operating under balanced-budget constraints and, since the 2025 reconciliation act, with reduced federal support.

Private-pay home care, assisted living, and hospital systems do bid wages upward, and that margin is real. But the administratively priced share is large enough to blunt the adjustment across the sector, so the shortage manifests less as wage inflation than as turnover, unfilled shifts, waiting lists, and the transfer of care onto unpaid family members, as documented in Section 3.4.

This is Loop C operating as a labor market phenomenon, and it produces a result that should trouble anyone forecasting from aggregate statistics. A sector with enormous excess demand will not generate the wage signal that normally accompanies it, because a plurality of its revenue is administratively fixed. **Analysts reading the aggregates will see a labor market that appears to be clearing while the underlying service is quietly rationed.**

4.3 What We Actually Know About AI and Jobs

Artificial intelligence comes with an evidence problem. The claims made loudest about it go well past what the data support, and the claims I can defend are narrower than what you hear. So let me separate what we know from what we are guessing, which matters more here than anywhere else in this paper.

What the aggregate data show is an absence. Studies of AI exposure and adoption in both the United States and Denmark find no detectable increase in aggregate unemployment among highly exposed workers since generative AI became widely available, and no significant effect on wages or hours (Humlum and Vestergaard 2025; Massenkoff and McCrory 2026). Mass white-collar unemployment has not arrived. Anyone claiming otherwise is arguing against the available evidence.

The disaggregated data show a concentrated effect at the entry point. Brynjolfsson, Chandar, and Chen (2025) work with payroll records covering roughly one in six American workers. They found that workers aged 22 to 25 in the most AI-exposed occupations saw a relative decline in employment of 13 to 16 percent after generative AI diffused, while workers aged 35 to 40 in the same occupations held steady or gained. Employment in entry-level jobs with low AI exposure did not decline. The pattern is specific: early-career workers, in exposed occupations, and not their more experienced colleagues.

That finding is contested on identification grounds. A Federal Reserve Bank of New York analysis of job postings finds that the vacancy divergence between AI-exposed and

insulated occupations began before late 2022, and that there is no divergence between junior and senior postings within exposed occupations. Both complicate the attribution of the entry-level slowdown to AI alone. Their reading is that AI is one factor among several rather than the principal driver (Audoly et al. 2026). Broader hiring has been weak since 2022, due in part to interest rates and post-pandemic normalization. The responsible summary is that a distinctive age-specific pattern exists in employment data, that AI is a plausible and partial cause, and that the effect size attributable to AI specifically is not yet settled.

As of mid-2026, the evidence base consists of early studies working with payroll records and vacancy postings over roughly three years. The findings are suggestive rather than definitive; large aggregate effects have not appeared, and a reader in 2030 will know things this section cannot.

4.4 What Happens When You Delete the Bottom Rung

Even at the low end of the plausible range, the entry-level mechanism deserves separate treatment, because its most serious consequence is not the one that shows up in an unemployment rate. Over a twenty-five-year horizon, it is more consequential than the immediate employment effect, and it interacts directly with the demographic contraction described in this paper.

Entry-level professional work transmits skill as well as providing employment. Junior lawyers learn judgment by drafting documents that a partner corrects. Junior analysts learn what a number means by building models that turn out to be wrong. Junior developers learn architecture by writing code that gets reviewed. The tasks through which people acquire tacit expertise are, with unfortunate precision, the codifiable tasks generative AI performs well. A firm optimizing this quarter has a clear incentive to have the model do the drafting and skip the correction.

If firms respond by eliminating the junior rung rather than redesigning it, the cost appears with a lag of roughly a decade, as a thinner pipeline of workers who have the experience to do senior work. **The harm is slow-acting and no quarterly statistic will capture it, and it arrives at the worst possible moment.** Section 4.1 established that the prime-age workforce stops growing in the 2030s. To the extent the cohorts entering that workforce have also been denied the apprenticeship function of entry-level work, the country faces a shortage of experienced workers layered on a shortage of workers.

This is a coordination problem rather than a technology problem, and it is worth naming as such. No individual firm captures the full return to training a junior worker, since the worker

may leave; thus, underinvestment in training is the expected outcome even when the aggregate return is high. AI lowers the immediate cost of skipping the investment. The mechanism is familiar from labor economics and does not require any assumption about AI capability beyond what is already observed. Section 5 shows the education system arriving at the same problem from the opposite direction, and Section 9 asks what could realign the incentive.

4.5 Why the Two Shocks Don't Cancel Out

Table 6 sets the two shocks side by side. The central observation is that they are close to orthogonal. AI pressure falls on codifiable cognitive work performed disproportionately by young, credentialed workers in metropolitan labor markets. The demographic shortage falls on work requiring physical presence, performed disproportionately by low-wage workers, immigrants, and unpaid family members, and is most acute where the population is oldest.

Table 6. The two labor shocks compared

	The AI shock	**The demographic shock**
Direction	May reduce demand for particular tasks and some categories of labor	Reduces supply of labor
Work affected	Codifiable cognitive tasks: drafting, coding, summarizing, routine analysis, entry-level professional work	Tasks requiring physical presence: direct care, nursing, food service, construction, skilled trades
Workers affected	Concentrated among the young and credentialed, especially ages 22-25 in exposed occupations	Employers in care and physical services, and the households that depend on them
Timing	Present tense, uneven, and contested in the data	Present tense in care, steepening sharply from the mid-2030s

Geography	Upper-branch metros, where AI-exposed occupations concentrate	Everywhere, but most acutely in older lower-branch counties
Distributional incidence	Ambiguous. Displaces upper-branch workers but accrues to upper-branch asset holders. Net effect depends on whether displaced workers hold equity	Raises wages where prices can adjust. Where reimbursement is administratively constrained it appears instead as vacancies, turnover, rationing, and reduced output, while raising costs for lower-branch households
Confidence	Contested; aggregate effects undetected, age-specific pattern documented but disputed	High; rests on cohorts already born and on observed sector vacancy data
Policy handle	Education, training, and the design of entry-level work	Compensation, immigration, credentialing capacity, and care technology

Source: author. The comparison is analytical rather than quantitative, and the two shocks differ in the confidence attached to them.

The tempting inference is that the two problems solve each other, that displaced analysts become home health aides and the market clears. They will not. The wage differential runs the wrong way. The work is physically and emotionally demanding in ways that do not attract those with alternatives. Credentialing takes time and money a displaced worker often lacks. The geography is wrong, with the surplus in metros and the acute shortage in aging non-metropolitan counties. And the move is one-directional in status terms, which people resist even at similar pay.

The mismatch is occupational and spatial at once, and the second half is the one usually left out. Health care administration illustrates both. A very large workforce is employed in health care administrative roles, management, claims processing, billing, coding, prior authorization review, and medical records, and the tasks are among the most codifiable in the economy. Those tasks are codifiable, concentrated in urban back offices and remote work hubs, and among the most exposed to automation of any large occupational category

in the country. The clinical and direct care shortages are dispersed and worst in rural counties.

Retraining a displaced billing specialist in a city as a rural nurse requires relocation as well as credentialing, and policy that funds the credential without addressing the move will leave the administrative workers in the metros and the shortages where they are.

The United States is heading toward a labor market with a surplus of one kind of worker and a shortage of another, separated by a gap that is occupational, geographic, and financial at once, and that markets alone do not bridge quickly. This describes a policy problem rather than a self-correcting adjustment. Section 9 takes up what could bridge it.

4.6 Scarcity Changes How Things Get Made

Standard economics predicts that labor scarcity induces offsetting adjustments rather than translating one for one into lost output. It changes how output is produced. Employers facing a shortage raise wages, substitute capital for labor, automate tasks, redesign processes, relocate activity, draw in workers from outside the labor force, or recruit from abroad. Section 2.1 discussed the participation margin. This subsection takes the capital margin, because an argument that productive capacity binds must say how much of the constraint the economy can engineer around.

The adjustment is real and highly uneven across sectors, and the unevenness is what makes the care argument work. Manufacturing, logistics, warehousing, agriculture, and back-office administrative work have high substitutability: the tasks are physically repetitive or informationally codifiable, capital equipment can replace hours of labor, and the historical record shows exactly this response to labor scarcity. Sections 4.3 and 4.4 describe its current version in cognitive work. In these sectors, a shrinking labor force is a problem the price system solves, at some cost and with distributional consequences, but it solves it.

Direct care, nursing, childcare, and construction tasks requiring presence have low substitutability because the reasons are structural rather than failure of engineering effort. Someone must be in the room to lift a person out of bed. Assistive technology helps at the margin, and Section 9.4 assesses how much. The measured productivity growth of the sector has been near zero for decades, which is what Baumol's analysis predicts for activities where the labor is the service rather than an input to producing it.

This asymmetry is the core of the paper's labor market argument, and it sharpens the claim rather than softening it. The economy will adapt to demographic scarcity in the sectors where adaptation is possible, and the adaptation will pull capital and attention toward those sectors because that is where the returns are. Care will absorb a rising share of national labor input precisely because it cannot adapt the same way**. A projection that applies an economy-wide productivity growth rate to a care sector defined by physical presence is not being optimistic. It is applying the wrong number to the wrong sector.**

4.7 Who Pays, and Where

The AI shock lands on upper-branch workers, which appears to cut against the framework. It does not. The gains accrue to capital, and capital ownership is concentrated at the top of the upper branch, so a displaced junior analyst and the shareholder capturing the productivity gain sit on opposite sides of the same transaction despite sharing a geography and an education. The young workers most exposed also hold the least accumulated capital, so an event that is modest in aggregate for their generation is severe for those without a cushion.

The demographic shock is simpler in its incidence and harsher. Rising care costs are paid disproportionately by lower-branch households, either in money or in foregone earnings when a family member leaves work to provide care.

Rising care wages, which would be a genuine gain for a low-wage workforce, are partly suppressed by the administrative pricing described in Section 4.2**. So the lower branch pays much of the cost of scarcity while receiving only part of the wage gain that scarcity would ordinarily produce.** That combination is a reasonably precise description of what the lower branch has experienced across multiple markets since 2020.

Geographically, the two shocks separate cleanly and reinforce the divergence. Upper-branch metros face an adjustment problem in the composition of professional work while retaining the in-migration that keeps their labor forces growing, and they retain the institutions, described in Section 2.8, that let them retrain and redeploy. Lower-branch counties face absolute labor scarcity in the sectors that keep an aging population alive, without the wage flexibility to attract workers and without the in-migration to replace them. Loop B accelerates: when the local hospital cannot staff a service line, it closes the line, and closure pushes out the working-age families who might have staffed it.

5. Education: The Pipeline Empties on Schedule

Education is where all this is most predictable. The students of 2044 have already been born. The students of 2032 are already sitting in classrooms. There is no uncertainty about how many there will be, only about what institutions do about it. That makes education the cleanest test we have of the argument in Section 1.

5.1 The Schools Are Already Emptying

The numbers are already in. Public K-12 enrollment peaked at 50.8 million in fall 2019 and had fallen to roughly 49.4 million by 2024, a loss of more than 1.4 million students. NCES projects a further decline to approximately 46.9 million by 2031, a 7.6 percent drop from the peak, with total enrollment lower in fall 2031 than in fall 2022 in 40 states and the District of Columbia, including all nine Northeastern states (NCES 2025a). The declines are steepest in elementary and middle grades, which is expected when the cause is births rather than migration: the shortage enters at the bottom and works upward one grade per year, arriving in high schools in the late 2020s and in colleges immediately thereafter.

The financial mechanism makes this difficult and is widely misunderstood. School costs are dominated by fixed and semi-fixed inputs: buildings, heating, transportation routes, administration, and a teacher in front of a class whose cost does not fall proportionally when the class has 18 students instead of 24. Revenue, by contrast, is substantially per-pupil, because most state aid formulas tie funding closely to student counts. A district losing 10 percent of its students loses close to 10 percent of its revenue and can shed nowhere near 10 percent of its costs without closing buildings.

The mechanism is easy to state and easy to underestimate, so Table 7 works it through for an illustrative district. The numbers are illustrative and rounded, but the structure is what matters: revenue falls with enrollment, and most costs do not.

Table 7. Why enrollment loss creates a structural deficit, illustrative district

	At 2,000 students	**At 1,800 students**	**What happens**
Revenue at $16,000 per pupil	$32.0 million	$28.8 million	Falls the full 10 percent, immediately
Buildings, utilities, maintenance	$4.8 million	$4.8 million	Unchanged. The building does not shrink

Transportation	$2.2 million	$2.1 million	Routes stay; buses run the same miles with fewer riders
Administration	$3.2 million	$3.1 million	Lumpy. You cannot cut a third of a principal
Instructional staff	$17.6 million	$16.4 million	Falls only if sections close; otherwise classes shrink
Other instructional and support	$4.2 million	$3.9 million	Partly variable
Total cost	$32.0 million	$30.3 million	Falls about 5 percent against a 10 percent revenue loss
Operating gap	balanced	$1.5 million deficit	About 5 percent of budget, in one year

Author's illustration. Per-pupil revenue is near the national average; the cost shares are representative rather than drawn from any particular district. The point is the asymmetry between the revenue line and the cost lines, which holds across a wide range of assumptions.

A five percent gap does not sound like a crisis. It is one, because it recurs every year the decline continues and because the options for closing it are all consequential: close a building, raise the local levy, defer maintenance, or cut programs. Most districts do some of each, and consolidation follows when the arithmetic no longer works. A district that loses ten percent of its enrollment over a decade is not managing a ten percent problem once. It is managing a compounding one.

Consolidation follows, and consolidation in a rural district means longer bus rides, the loss of the school as a community institution, and frequently the departure of the families with the means to leave. That is Loop B running through the school system, and it is the most common way the loop is actually experienced by the people inside it.

State policy mediates how sharply this lands. Funding formulas differ, and a minority of states use hold-harmless provisions, enrollment averaging, or minimum district grants that slow the revenue decline and give districts time to adjust their cost base. Others pass the loss through in a single year. The same demographic event therefore produces a manageable adjustment in one state and a fiscal emergency across the border, which is a reminder that the institutional design problem in this paper is substantially a state-level problem and that some of it is already solved somewhere.

The distributional pattern here is not a simple upper-branch and lower-branch split. Some of the sharpest recent declines have occurred in high-cost coastal metros where housing prices have driven out families with children, and in high-income districts where families shifted to private schooling and homeschooling after 2020. This is the pattern Section 2.8 anticipated. What distinguishes the lower branch is not the rate of decline but the absence of alternatives. A wealthy suburban district that loses enrollment consolidates two elementary schools and continues. A rural county that loses enrollment closes the only school within thirty miles, and the closure is functionally permanent because nothing will bring the families back.

5.2 Fewer Students Won't Fix the Teacher Shortage

You might expect falling enrollment to solve the teacher shortage that has preoccupied American education policy for two decades. It will not, and the reason illustrates the mismatch logic of Section 4 operating inside a single sector.

Teacher shortages have never been general. They are specific to subjects and to places: special education, mathematics, the sciences, career and technical fields, and instruction for English learners, concentrated in rural districts and high-poverty urban schools. Falling enrollment reduces the aggregate number of teaching positions, but it does not reallocate a surplus elementary generalist in a shrinking suburb into a vacant rural special education position four hundred miles away. Aggregate contraction and specific shortage coexist for the same reason surplus analysts and scarce home health aides coexist. The categories do not convert.

Contraction also makes the specific shortages harder to fill. A district losing enrollment and revenue has less capacity to pay a premium for a hard-to-staff subject, and consolidation lengthens commutes and reduces the number of positions in any given community, which shrinks the pool of candidates willing to live there. Small rural districts often cannot fill a specialized position at all and cover it through long-term substitutes or

online instruction. The teacher workforce is also aging alongside the population it serves, which means retirement-driven vacancies continue arriving even as total positions decline.

And the pipeline is thinning at the source: teacher preparation enrollment has fallen substantially over the past fifteen years, and the institutions that supply most teachers are the regional public universities and small private colleges under the greatest financial stress, which is the subject of Section 5.3.

The result is a sector that will appear overstaffed in the aggregate while remaining unable to deliver instruction in specific places. Policy responses calibrated to the aggregate, hiring freezes and position eliminations, will worsen the specific shortages, because the positions easiest to cut are rarely the ones hardest to fill.

5.3 The Enrollment Cliff Is Here

The higher education version arrives on a date that has been known for eighteen years, because it was fixed by the birth data of 2008. The number of American high school graduates peaked in 2025 at roughly 3.8 to 3.9 million and is projected to fall about 13 percent through 2041, removing roughly 576,000 graduates from the annual pipeline (WICHE 2024). Grawe (2018), whose projections have held up well, estimates a decline in the college-age population of roughly 12 to 15 percent between 2025 and 2029 across most regions. Fall 2026 is the first cohort of the decline. The cliff is not approaching. It is here.

And three factors compound the demographic one. College-going rates among high school graduates have fallen, from about 70 percent in 2016 to the low 60s, driven by cost, skepticism about return on investment, and the relative appeal of skilled trades. International enrollment, which many institutions used to offset domestic softness, has contracted under the same visa and immigration policies analyzed in Section 2.4, with new international enrollments falling sharply in fall 2025 and graduate programs hit hardest (IIE 2025). And the pandemic pulled forward part of the decline. Institutions face a shrinking pool, a falling capture rate within that pool, and the loss of their principal hedge, all at once.

Financial consequences follow quickly for tuition-dependent institutions with small endowments. Fitch assigned the sector a deteriorating outlook for 2026, the second consecutive year, and more than 100 colleges, predominantly small private nonprofits, carry elevated risk of closure or merger (Fitch Ratings 2025). Closures have been running at roughly a dozen or more nonprofit institutions a year and are widely expected to accelerate through the late 2020s. The geography of closure risk maps onto the geography of Section

2.5 with some precision: small private colleges in the rural Northeast, upper Midwest, and Rust Belt, institutions located in counties already losing population and serving regional student bodies shrinking faster than the national average.

Two consequences deserve emphasis beyond the institutional accounting. First, a college closure in a small city is a major employer closure, often the largest, in a place with no replacement. The higher education contraction is therefore also a rural and small-city economic shock, landing in the same counties absorbing hospital service line closures and school consolidations.

Second, public flagship and highly selective private institutions will be largely insulated, since they draw nationally and can maintain enrollment by moving down their applicant pools. The contraction will thus concentrate access loss where alternatives are thinnest, widening the educational component of the K-shaped divergence rather than distributing the pain evenly.

5.4 Closing the Schools We Are About to Need

The sharpest irony here is sequencing, and it is Loop D of Figure 1. Sections 3 and 4 established that the country needs to recruit, retain, and replace workers across several million projected direct care openings, expand nursing and allied health pipelines, and retrain workers displaced from AI-exposed occupations into work AI cannot do. Essentially all of that capacity sits in community colleges, regional public universities, and the small private institutions now under the most financial stress. **The education system is contracting on a demographic schedule at precisely the moment the labor market needs it to expand in specific directions.**

These are not always the same institutions, and community colleges have grown recently on dual enrollment and affordability. But the sector-wide financial stress, the closure risk concentrated in non-metropolitan areas, and the consolidation of regional public campuses all cut against training capacity where the care shortage will be most acute. An aging rural county that loses its community college nursing program does not regain it when demand rises, because the program closed for enrollment reasons unrelated to the demand that materializes five years later. The loop closes regionally. Fewer trained workers means thinner local services, which pushes out the young families whose children would have refilled the enrollment that emptied in the first place.

There is a corresponding curricular problem that Section 4.4 anticipated. If generative AI is displacing entry-level professional tasks, the traditional four-year credential aimed at those roles delivers a weakening return exactly as its price becomes harder to justify,

which is part of why the college-going rate is already falling. Institutions face a contraction in the number of students and a simultaneous question about what the product should be. Handling either well is difficult. Handling both at once, under financial stress, with governance structures designed for growth, is the actual challenge facing American higher education for the next fifteen years.

5.5 Where That Leaves Us

Labor force growth slows sharply over the projection period and stops under low immigration, so trend output growth must come almost entirely from productivity, hours, quality, or capital intensity. Fiscal projections assuming historical growth alongside these demographic assumptions are embedding a demanding and usually unstated assumption. Second, AI and demography are two shocks running in opposite directions on different workers, and the surplus one creates does not fill the shortage the other creates. Third, the care sector cannot fully clear through wages because a plurality of its revenue is administratively priced, so shortage will present as rationing rather than as measurable wage inflation.

Fourth, the education system contracts on a fixed and now-known schedule, the contraction concentrates in the same counties absorbing every other adjustment in this paper, and the retraining capacity the labor market will need sits precisely in the institutions under the most stress. Section 6 turns to the political system, where these accumulated pressures meet the reapportionment cycle.

Table 8. Labor market and education dates established in these sections

Date	Milestone
2025	U.S. high school graduates peak at roughly 3.8 to 3.9 million
2026	The higher education enrollment cliff begins
2025-2029	The college-age population falls roughly 12 to 15 percent
By 2029	Institutional closures accelerate; 100 or more colleges at elevated risk
2031	K-12 enrollment falls to roughly 46.9 million, 7.6 percent below the 2019 peak

2034	Labor force participation reaches 61.1 percent, about 4.3 million fewer participants
2024-2034	Direct care records 9.7 million openings, more than any other category
2041	High school graduates fall 13 percent from the 2025 peak

Education entries rest on cohorts already born and carry unusually high confidence. Confidence classes and sources appear in the consolidated timeline, Table 12 in Section 10.

6. How Population Turns Into Political Power

The American political system turns population into power through a machine that runs once a decade and cares about nothing in this paper except headcount. The census counts people. Apportionment turns counts into House seats and electoral votes. Redistricting turns seats into districts. Every step is mechanical, and every step is fought over. Two full cycles fall inside the window of this paper: 2030 and 2040. My paper on the 2030 census (Lane 2026b) treats the redistricting law and strategy in detail; the purpose here is narrower, to establish the demographic transmission and its interaction with the fiscal deadlines of Section 3.

6.1 The Reapportionment Machine

Reapportionment is the purest expression of the sorting described in Section 2.5, and a clean illustration of the growth-institution problem from Section 1. Because the House is fixed at 435 seats, apportionment is zero-sum: a state gains representation only by growing faster than the national average, and states can lose seats while gaining population.

Apportionment is the constitutional instance of a dynamic that now operates throughout American federalism, and Section 9.2 treats the general case. In an era of national near-stagnation, that mechanism converts modest differences in growth into transfers of political power. Published projections based on 2024 and 2025 population estimates converge on the same pattern, with Texas and Florida the principal gainers and California, New York, and Illinois the principal losers.

Table 9. Projected 2030 reapportionment, published estimates

State	Projected seat change	Range across published estimates	Demographic driver
Texas	Gain 3 to 4	+2 to +4	Domestic in-migration, comparatively high fertility, large immigrant population
Florida	Gain 2 to 3	+2 to +4	Retiree and working-age in-migration, offset by affordability-driven slowing
Georgia, North Carolina, Arizona, Idaho, Utah	Gain 1 each	0 to +1	Metro growth and domestic in-migration
California	Lose 2 to 4	-2 to -4	Domestic out-migration on housing cost, reduced international migration
New York	Lose 1 to 3	-1 to -3	Sustained domestic out-migration; delegation continues an eight-decade decline
Illinois	Lose 1 to 2	-1 to -2	Domestic out-migration, natural decrease in downstate counties
Minnesota, Oregon, Rhode Island, Pennsylvania	Lose 1 each	0 to -1	Slow growth against faster-growing states

Compiled from Brennan Center (2026), Election Data Services (2025), and Cervas (2026). Estimates vary because they extrapolate from different base years, and early-decade projections have historically been off by a seat or more in several states.

Two cautions belong with these numbers. The first is projection risk. Reapportionment estimates made five years before a census extrapolate migration flows that respond to housing costs, interest rates, and labor markets. Florida's projected gain has already fallen from four seats to three as affordability slowed in-migration, and one recent year of growth alone would imply only two. The same mechanism could slow Texas.

The second is enumeration risk. Apportionment depends on the accuracy of the count, not on the underlying population, and the 2020 census produced measurable undercounts and overcounts that varied by state and demographic group. A count conducted under fiscal or political pressure, or in which immigrant households are reluctant to respond, will allocate seats differently than the true population would. Enumeration accuracy is policy-contingent rather than demographic.

Composition matters more than the headline count. The gaining states are not uniformly young, and treating Sun Belt growth as one phenomenon obscures the mechanism. Florida is the country's principal retiree destination, so seats moving there move toward an older median population with unusually high dependence on Social Security and Medicare.

Texas gains from a different mix: domestic in-migration, comparatively high fertility, and a large immigrant population. Under the low-immigration scenario of Section 2.4, Texas gains fewer seats and the entire Sun Belt shift is muted, which is the clearest illustration in this paper that immigration policy and political representation are the same subject viewed from different angles. Restriction reduces the population that produces representation in the states whose delegations have been growing.

6.2 What Callais Changed

Section 2.6 established that the country becomes majority-minority around the mid-2040s, with the crossover already past among people under 18. Whether that compositional change converts into proportional political representation depends on redistricting law, and that law changed materially in April 2026.

In Louisiana v. Callais, decided 6-3 on April 29, 2026, the Supreme Court held that Louisiana engaged in an unconstitutional racial gerrymander when it drew a second majority-Black congressional district to comply with Section 2 of the Voting Rights Act. The Court narrowed the circumstances in which a racial vote dilution claim can succeed,

raising the evidentiary burden on plaintiffs and permitting states to defend maps by asserting partisan rather than racial motives. It did not formally overrule Section 2.

In dissent, Justice Kagan wrote that the decision renders the provision close to a dead letter, and voting rights organizations have characterized the ruling as eviscerating it. Defenders argue that it restores the Equal Protection Clause's general prohibition on race-conscious districting and resolves a tension between Section 2 compliance and the Fourteenth Amendment that the Court had acknowledged for years without settling. Both characterizations describe the same operative change: race-conscious districting to remedy vote dilution is now harder to require and easier to challenge.

The demographic significance holds under either characterization. For six decades, one mechanism by which population changes translated into representational change was litigation under Section 2. That channel is now substantially narrowed, immediately before the first redistricting cycle in which the demographic shifts of this paper become visible in the data. The states gaining seats under Table 9 are gaining them largely because of the growth of nonwhite populations, particularly in Texas, Georgia, North Carolina, and Arizona, and those are precisely the states where Section 2 litigation had previously shaped district lines.

Population growth will produce seats. Whether it produces representation for the populations that generated the growth is a more open question than it was in 2025, and the answer will be set by state legislatures rather than by courts.

The general point survives the particular case, which matters for a paper like this one with a 2050 horizon. Doctrine may evolve again before the 2031 cycle, and this analysis does not rest on Callais remaining good law. What it rests on is the structural observation: population change does not automatically become representation. Census enumeration, apportionment arithmetic, redistricting institutions, and voting rights doctrine each mediate the translation, and each is a policy variable. Anyone reasoning from the mid-2040s majority-minority crossover to a conclusion about political outcomes is skipping the step where composition becomes representation, and that step has always been mediated by law. It is outside the scope of this paper to take a position on how legislatures should draw maps.

6.3 Having the Numbers Isn't Having the Votes

Here is where it gets interesting. The age structure of the electorate diverges from the age structure of the population, and the gap runs in a consistent direction. In the 2024 presidential election, Census data show turnout of roughly 75 percent among citizens 65

and older, against roughly 48 percent among those 18 to 29, with overall turnout at 65.3 percent (U.S. Census Bureau 2025). Older Americans therefore hold a share of the actual electorate meaningfully larger than their share of the population, and the gap is wider in midterm and primary elections.

This produces a lag of specific and consequential length. Millennials and Generation Z reach a majority of the eligible electorate by headcount in the mid-2030s. Weighted by turnout, they reach a majority of actual voters some years later. The interval between those two dates is the period in which the largest distributional questions in this paper come due: the forecasted 2032 Social Security deadline, the 2033 Medicare deadline, and the Medicaid long-term care pressure that steepens from the mid-2030s. The decisions get made while the older cohort still holds the effective electoral majority, on programs whose costs fall on the younger cohort.

This is not a claim that older voters will vote selfishly. Polling on entitlement questions does not support a simple generational split, and cohort attitudes are not monolithic on any of these questions. It is a claim about the electoral environment in which decisions get made, shaped by turnout rather than population shares.

6.4 Everything Lands at Once

Three institutional clocks converge in a narrow band, and Figure 2 sets them out. The demographic clock reaches its own milestones around 2030, as the last boomers turn 65 and deaths begin to exceed births. The political clock runs the 2030 census, the 2031-2032 redistricting cycle, and a presidential election in November 2032 under a redrawn map. The fiscal clock reaches OASI depletion in the fourth quarter of 2032, within weeks of that election, and Medicare HI depletion in mid-2033.

Figure 2. Three institutional clocks converge, 2030-2033

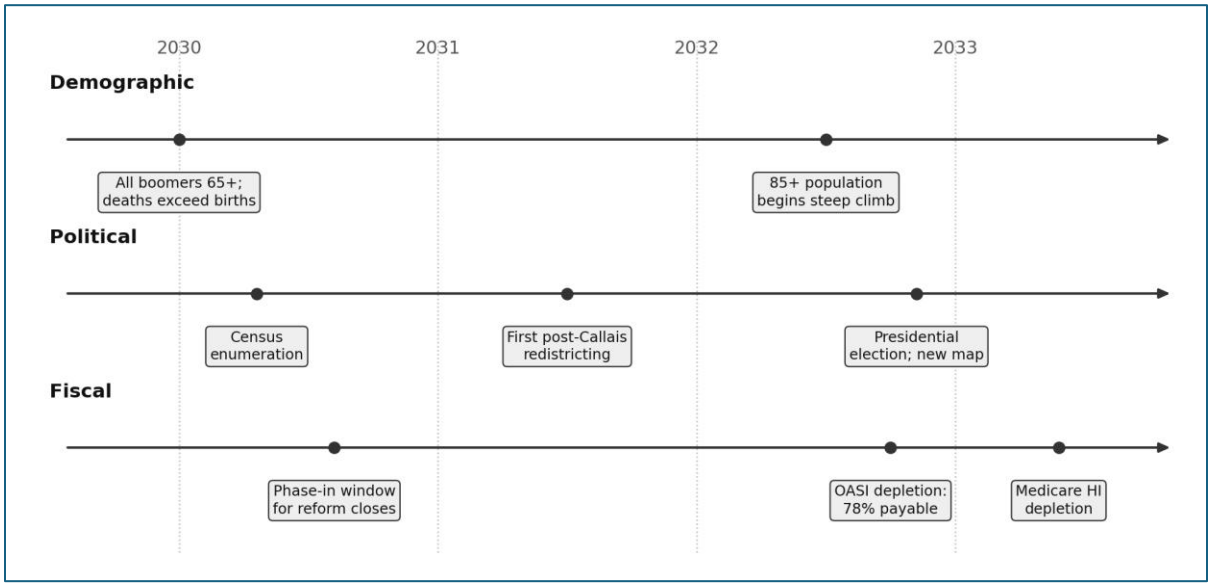

Source: author. Census and apportionment dates are fixed by statute. Trust fund dates reflect the 2026 Trustees Reports under current law and assume no legislative action.

These calendars arise independently. The census schedule is constitutional, the presidential election calendar has nothing to do with the baby boom, and the trust fund dates follow from program design and actuarial revision.

What demography supplies is not the timing but the significance: because the boomer cohort is passing through the age structure, the fiscal dates fall where they do, and because internal migration has been persistent and directional, the reapportionment is consequential. The convergence is a coincidence of calendars given weight by a demographic transition, and the practical result is that **the largest scheduled fiscal decisions in the American social contract will be litigated in a presidential campaign, on a map drawn under contested rules, within a twenty-four-month window.**

6.5 The Political Paradox

Section 3.2 established that a benefit cut would fall hardest on counties where federal old-age transfers are the largest source of personal income, and that those counties are concentrated in the lower branch of the K. The same counties have increasingly supported the coalition that advocates for limiting future federal retirement and health benefits. This support is evident in think tank agendas, committee proposals, and budget resolutions, but not in presidential campaign rhetoric, where both parties typically pledge to protect

these benefits. I have called this the Political Paradox in earlier work (Lane 2026a). It needs careful statement, because the careless version is condescending and wrong.

The careless version says voters are acting against their interests out of ignorance. Several better explanations are available. Social Security and Medicare are not perceived as redistribution by their recipients; they are perceived as earned insurance, paid for through a lifetime of payroll deductions, and that perception is largely accurate as a description of how the programs were designed and sold. Voters therefore do not read a candidate's general position on federal spending as a threat to a benefit they regard as their own property. Second, no major candidate campaigns on cutting Social Security, so the issue is not on the ballot in a form that would produce the vote the paradox implies.

Third, **the issues that do drive these electorates, cultural questions, immigration, and the perceived indifference of national institutions to their communities, are real preferences rather than errors.** The affordability grievance running underneath them is precisely what the K-shaped framework predicts: these are the counties that have experienced the sharpest erosion of local opportunity, and party competition is more salient on that erosion than on transfers nobody is currently proposing to cut. All of these explanations operate on averages. Individual voters hold varied combinations of views about redistribution, cultural questions, and program design, and nothing here should be read as a characterization of any particular county or group.

A fourth explanation subsumes the other three, and it is the one I find most persuasive. Populism, in Cas Mudde's formulation, is a thin-centered ideology (Mudde and Rovira Kaltwasser 2017): it supplies a division between ordinary people and a self-serving elite and attaches that division to whatever host ideology is locally available. It arises reliably from the combination of economic dislocation, political exclusion, and cultural alienation.

American history offers four distinct episodes of it since the 1890s, arriving from the left as often as from the right. What changes between episodes is not the structure of the grievance but the identity assigned to the elite, which has been in turn financiers, subversives, bureaucrats, and cosmopolitan insiders. I trace that sequence in earlier work (Lane 2025).

Read this way, the lower-branch voter is not misreading his interests. He is registering a grievance that is accurate on its own terms, about the erosion of local opportunity documented throughout this paper, through the only vehicle available. Neither party has presented an economic program addressed to county-level decline, so the grievance attaches to the cultural and institutional questions that are on the ballot. This is a rational reaction to a limited menu, not a comprehension failure, and it suggests that grievance will

align with whichever framing is present when the benefit question emerges. This framing is not fixed beforehand.

The paradox establishes structural fragility rather than voter failing. The arrangement depends on the benefit question never becoming salient. In 2032, it becomes salient by operation of law, on a date fixed in advance, with an automatic 22 percent reduction as the default outcome. The political economy that has held since roughly 2016 has not been tested against some concrete, dated, unavoidable benefit cut landing on the most loyal counties of the governing coalition. Section 9 returns to this, because it bears on which policy responses are actually available: the feasible set is shaped less by what is fiscally sensible than by which coalitions can survive the vote.

6.6 Where That Leaves Us

The demographic shift transfers House seats and electoral votes toward the South and Mountain West in 2030 and again in 2040, but the magnitude is uncertain, depends on migration flows already slowing and on immigration policy, since the low scenario mutes the entire shift.

Second, the conversion of population change into representation is mediated by enumeration, apportionment arithmetic, redistricting institutions, and voting rights doctrine, each a policy variable, so reasoning directly from the majority-minority crossover to political outcomes skips the step that does the work.

Third, the political and fiscal calendars converge inside twenty-four months in the early 2030s, before an electorate still weighted toward the cohort receiving benefits rather than the one financing them. Before turning to the household, Section 7 examines the level of government that must absorb many of these pressures without the federal government's monetary latitude.

7. The Governments That Cannot Print Money

Everything so far has treated the fiscal problem as a federal one. Section 3.1 argued that federal trust fund dates are legal triggers rather than measures of national capacity, because the federal government issues the currency its obligations are denominated in. That argument is correct, and it has a boundary that matters as much as the argument itself.

States and municipalities cannot do any of it. They face hard balanced-budget requirements. They borrow in a currency they cannot create. Their obligations are

enforceable against a tax base that demography is reshaping under them. The functional finance argument that applies to Washington emphatically does not apply to Topeka, and for subnational government the financing constraint is real in exactly the way Section 3.1 says it is not for the federal government.

That asymmetry is the most important single fact about which level of American government can absorb this transition, and it cuts the wrong way. The level least able to run deficits is the level that operates the schools, hospitals, water systems, roads, and emergency services, which are precisely the fixed-cost activities that demographic contraction makes expensive per capita. Federal choices then transmit downward. The 2025 reconciliation act reduces federal Medicaid spending by an estimated $911 billion over ten years and constrains the provider taxes states have used to finance their share (KFF 2025). The federal saving is realized as a state obligation.

7.1 What Happens to the Tax Base

Aging by itself does less damage to state revenue than one might expect. Pew's 2026 review finds the effect modest in many states once age composition is separated from everything else happening alongside it. The pressure shows up on the spending side, where the evidence is stronger (Pew 2026).

What creates real fiscal risk is the combination: aging, plus outright population loss, plus weak property values, plus working-age taxpayers leaving, plus rising service demand. It is the coincidence that does the damage, not the age structure by itself.

Within that qualification, several mechanisms run the same direction. Retirees consume less of what is taxable, and in most states they receive some exemption on retirement income or property tax relief on age grounds. Those exemptions were designed when the elderly were a small share of the community. They become expensive when the elderly are the plurality, and they are politically difficult to withdraw from the group with the highest turnout, for the reasons Section 6.3 describes. Pew documents that states have begun modeling this directly, with Colorado and Pennsylvania incorporating aging into long-term budget assessments and Montana's 2026 study projecting slower income tax revenue growth, rising healthcare needs, and declining labor force participation together.

Property values in shrinking counties stagnate or fall, which caps the base directly. And because assessed value responds to transaction activity, a market with few sales is slow to register decline, so the fiscal shock arrives late and then arrives at once. Jurisdictions that look stable on assessed value can be several years behind their own market.

7.2 The Costs That Do Not Shrink

This is the mechanism Section 5.1 describes in school districts, generalized to everything a local government owns. A water system sized for 40,000 people does not become cheaper when 30,000 remain. The pipes are in the ground and the treatment plant is built. The same holds for sewer networks, road miles, bridges, emergency dispatch, fire coverage, and public buildings.

Per capita costs therefore rise mechanically as population falls, and a shrinking jurisdiction must charge more per household to deliver the same service. Deferred maintenance is the usual response, and it is a loan taken against a future in which fewer people will be available to repay it. The obligation does not disappear. It converts from an operating cost into a capital liability, and it is not recorded anywhere in the annual budget that made the choice.

Two independent literatures confirm the mechanism and establish its scale. The OECD finds across member countries that low-density development and outdated spatial plans in shrinking regions produce underused infrastructure and rising per capita service costs, and that the cost of maintaining essential infrastructure barely varies with population because fixed costs dominate (OECD 2025b). Sutradhar, Spearing, and Derrible (2024) put American numbers on the exposure, projecting that close to half of the roughly 30,000 U.S. cities face some population decline by 2100, affecting 27 to 44 percent of populated area, with implications for water, electricity, transit, and internet provision.

7.3 Pensions and Other Promises

State and local pensions are not a subnational Social Security, and treating them that way gets the mechanism backwards. The federal program is pay-as-you-go, so its finances depend directly on the ratio of workers to beneficiaries. State and local plans are funded systems holding real assets, roughly $6.5 trillion across the universe of them. Their condition depends on actuarial funding policy, contribution rates, investment returns, assumed discount rates, benefit design, and the existing funded ratio. A well-funded plan does not become insolvent because the sponsoring jurisdiction loses population.

The correct statement is narrower but still consequential. Demographic contraction does not make a funded pension insolvent; it magnifies the burden of an underfunded one. Required employer contributions and legacy liabilities are paid from current revenue, and that revenue rests on a payroll and tax base that population loss shrinks. A plan at 60 percent funding in a growing county and the same plan in a county losing a fifth of its

population over two decades are the same actuarial problem carried by very different capacity to pay. The demographic variable operates on the denominator, not the liability.

Two features then make the obligation harder than the federal analog rather than easier. It is frequently protected by state constitutional provision or contract law and cannot be reduced for current retirees in many jurisdictions, so the repair options available to Congress are unavailable to a governor. And it must be paid in a currency the state cannot issue. Where Section 3.1 argues that the federal trust fund date is a deadline for political decision, an underfunded subnational pension in a shrinking jurisdiction is closer to an ordinary debt with a deteriorating revenue pledge behind it.

7.4 What Shows Up in the Credit Markets

Persistent population decline eventually reaches debt service capacity, and the transmission runs through several credit sectors that share the same demographic driver. General obligation debt of shrinking municipalities weakens as the tax base ages. Revenue bonds of water and sewer systems weaken as connection counts fall while fixed plant remains. Rural hospital debt weakens where the payer mix is worst and the population oldest, which Section 3.4 documents alongside roughly 136 rural hospital closures since 2010. Higher education debt weakens at the tuition-dependent institutions facing the 13 percent decline in high school graduates through 2041 described in Section 5.3.

These exposures are typically analyzed as separate sectors by separate analysts using separate frameworks. They share a common geographic demographic risk factor, and they may deteriorate together in the same counties in the same decade. They are not a single exposure: their revenue sources, reserves, rate-setting authority, state support, legal protections, and essentiality all differ. The claim advanced here is that apparently diversified holdings may carry correlated demographic risk, which is a hypothesis the screen in Appendix A is designed to test. Appendix A develops the implication for anyone holding them.

7.5 Why Some States Can Absorb This and Others Cannot

A state with a growing, diversified, asset-heavy tax base can cross-subsidize a declining county indefinitely, and several do. A state where most counties are declining at once has no internal cushion, and its cities compete for a shrinking pool of taxpayers under the dynamic Section 9.2 describes. This is the fiscal expression of absorptive capacity as Section 2.8 defines it, and it is why the same national demographic trend will produce municipal distress in some states and a manageable reallocation in others.

State policy also mediates how sharply contraction lands, which means part of this problem is already solved somewhere. School funding formulas differ, and a minority of states use hold-harmless provisions, enrollment averaging, or minimum district grants that slow the revenue decline and give districts time to adjust their cost base. Others pass the loss through in a single year. The same demographic event produces a manageable adjustment in one state and a fiscal emergency across the border. The institutional design problem in this paper is substantially a state-level problem, and the variation is a resource rather than only a complication.

7.6 What Managing Decline Actually Looks Like

The preceding subsections describe what goes wrong. They do not answer the question a state budget director actually faces, which is what a rational government does when population decline is permanent rather than cyclical. That question has been almost absent from American policy discussion, because the assumption that decline is temporary and reversible is embedded so deeply in local planning that its alternative is rarely stated.

The alternative has a name. The OECD framework for depopulating regions distinguishes mitigation, meaning strategies to attract and retain population, from adaptation, meaning strategies to sustain quality of life at a smaller population (OECD 2025b). The distinction matters because the two call for opposite investments. Mitigation builds capacity in the hope that people arrive. Adaptation right-sizes capacity to the people who remain. American jurisdictions overwhelmingly pursue mitigation, often through tax abatements and business attraction incentives, and the OECD assessment of similar programs across member countries is that they rarely reverse a sustained demographic trend and frequently consume revenue that adaptation would have used better. Section 9.2 explains why they persist anyway.

Adaptation is politically brutal, which is the real obstacle. Its instruments are well known and each has a constituency opposed to it: decommissioning infrastructure rather than maintaining it, consolidating school districts and municipal service departments, regionalizing functions such as emergency dispatch and water treatment across jurisdictions that have historically guarded them, land banking and demolishing vacant property, and concentrating investment in a smaller number of viable service centers rather than distributing it thinly across every settlement that once had a population. Each amounts to a public decision that a place will be smaller permanently, which is the one thing local officials are least able to say.

The fiscal case for saying it is nonetheless strong. Every year a jurisdiction maintains a service footprint sized for a population it no longer has, it transfers cost onto a smaller base and defers maintenance it will eventually capitalize. The OECD recommendation is to adapt spatial planning and infrastructure to smaller populations rather than treating renewed growth as the only acceptable outcome, and to invest in access to services rather than in the physical footprint that historically delivered them, which is where telehealth, mobile service delivery, and regional consolidation enter. None of this is costless and none of it restores what was lost. The comparison is not against the county as it was in 1990. It is against the same county in 2045 having spent twenty years financing an unaltered footprint.

International experience suggests one further point, and it is the one most relevant to the framework of this paper. Where national population growth stops, economic geography tends to concentrate rather than shrink proportionally. Japan, furthest along this path among large rich countries, has seen continued consolidation around fewer larger urban centers alongside deeper rural depopulation, and the OECD finds the same pattern across member states, with more than a third of remote regions losing population since 2001 while urbanization continued. The American reading is not that the United States becomes Japan. It is that a national demographic plateau does not produce uniform national decline. It produces sorting, which is what Section 2.5 documents at the county level and what this section treats as the driver of fiscal divergence.

One observation belongs at the end of this section because it unifies three others. The American fiscal system has a hierarchy of absorptive capacity, and it runs in a consistent direction. The federal government can change financing rules and faces a real-resource constraint rather than a financial one. A state faces a balanced-budget requirement and borrows in a currency it cannot issue. A municipality has a water system it cannot resize. A household has almost no risk-pooling capacity at all.

Demographic burdens move downward through exactly that hierarchy, and each step down reduces the capacity to absorb them. A reduction in federal Medicaid spending becomes a state budget problem, which becomes a provider reimbursement problem, which becomes an unpaid family caregiver. Federal choices about training and education become regional institutional closures, which become local workforce shortages. At every transfer the recipient has less flexibility than the party that passed it along.

Stated plainly: the American fiscal system pushes demographic adjustment toward the levels of government, and ultimately the households, with the least capacity to absorb it.

The mechanism connects Section 3, this section, and Section 8, and it is a design feature rather than an accident of any particular policy.

Figure 3. The hierarchy of demographic adjustment

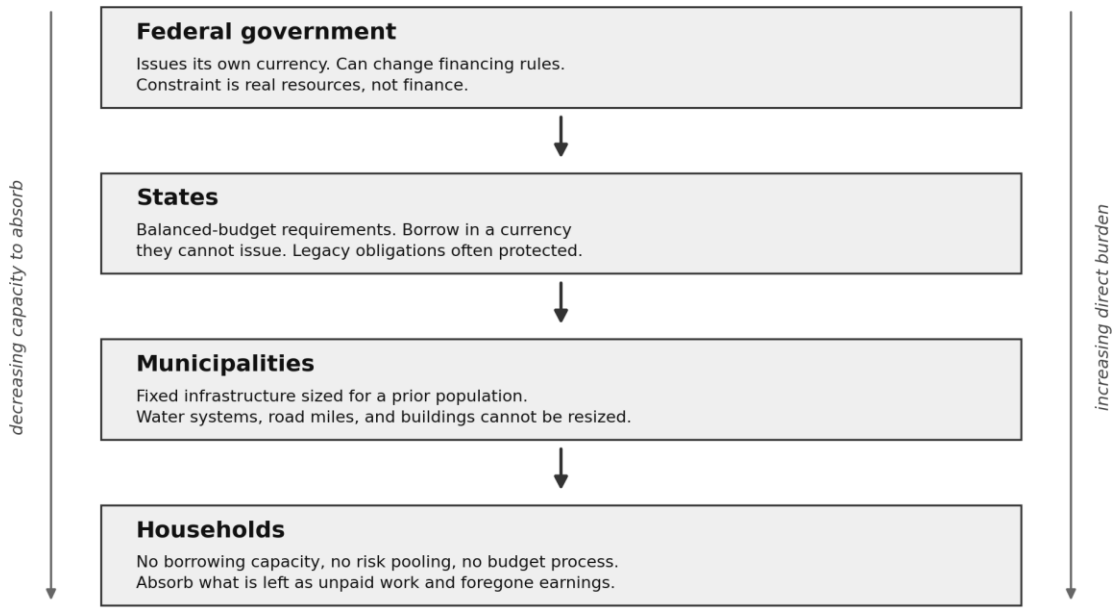

Source: author. Each level can transfer demographic cost to the level below it, and each recipient has less capacity to absorb that cost than the institution passing it along. Developed in Sections 3.1, 7, and 8.

8. What This Looks Like at the Kitchen Table

Everything in the preceding sections eventually arrives at somebody's kitchen table. When a trust fund is depleted, a household absorbs a smaller check. When a state trims home and community-based services, a daughter reduces her hours. When a rural hospital closes a service line, a family drives ninety minutes or does without. This section follows the adjustment into the household, which has no budget process, no borrowing capacity, and no way to register the burden it absorbs. Section 7 treated the parallel problem at the level of state and local government, where the same demographic pressure meets balanced-budget requirements, fixed infrastructure, and legacy pension obligations.

8.1 The Wrong Houses in the Wrong Places

By 2035, there will be roughly 50 million households headed by someone 65 or older, up nearly 20 million in two decades (JCHS 2016). Their housing does not fit them. Only a small share of American homes has basic accessibility features, such as a no-step entry with a bedroom and full bath on the entry level, and most older adults live in single-family houses they own and intend to keep. Even after age 80, only about one in twelve older adults lives in a nursing home, assisted living facility, or similar setting. The other eleven are in ordinary houses.

The dominant pattern for the 2030s and 2040s is therefore aging in place in housing not designed for it, converting a demographic fact into retrofit demand at a scale the construction sector has not planned for. Older homeowners already account for nearly half of home improvement spending, up from a historical 30 percent, though little of it targets accessibility.

The mismatch runs in both directions, and the second direction is the one that connects to fertility. Older households tend to live in family-sized homes within established neighborhoods that have good schools. They are remaining in these homes longer than earlier generations, due to preferences, a lack of attractive smaller options in the same area, and, since 2022, mortgage rates that are significantly lower than those for new purchases. That reduces the turnover through which family-sized housing historically reached the next generation of families.

Young households facing high prices and low inventory delay formation, and Section 2.2 established what delayed household formation does to completed fertility. The housing market is thus a transmission belt between the two ends of the age distribution: the housing decisions of the old constrain the family formation of the young, which determines the size of the cohort that will eventually support both.

The K-shaped incidence is straightforward. An upper-branch household ages in place in an appreciated asset, finances retrofitting from equity, and retains options including a home equity line, a reverse mortgage, or a sale into a strong market.

A lower-branch household ages in place in a home worth less than the cost of the modifications it needs, in a county where the market is thin and a sale may not clear. Roughly one in five older homeowners in the rural South live in a manufactured home, a housing type that is difficult to modify and depreciates rather than appreciates. The same demographic event, aging in place, is an asset management question on one branch and a habitability question on the other.

8.2 Most Care Is Unpaid, and Nobody Counts It

So far, the care economy has appeared in this paper as a labor market, a Medicaid line item, and a productivity problem. It is also, and primarily, a household activity. Most care in the United States is unpaid, performed by relatives, and invisible to every statistic that governs policy. Section 3.4 gave the scale.

In 2024, 59 million Americans provided care to an adult, delivering 49.5 billion hours worth roughly $1.01 trillion at replacement wages. That is equivalent to nearly 24 million full-time workers, about 17 percent of the full-time workforce, and it exceeds total federal, state, and local Medicaid spending that year (AARP 2026). The number of caregivers has risen sharply, up roughly 45 percent since 2015 (AARP and NAC 2025).

Unpaid family care is already worth more than everything federal, state, and local government spends on Medicaid combined. It is the larger part of the system, and the paid sector is the supplement.

This is not a supplement to the formal care system. It is the larger part of the system, and the formal sector is the supplement.

Two features of this arrangement determine how the coming decades go. The first is that it is the residual claimant. Every constraint elsewhere in the system, a Medicaid waiver waiting list, an unfilled home health shift, a nursing facility that has closed, a premium that consumes a fifth of a Social Security check, resolves by transferring work to an unpaid family member. That is why the system appears to clear. Demand does not disappear when supply is rationed; it moves to the household, where it is not counted. Any projection of long-term care costs that does not state its assumption about family absorption is implicitly assuming the residual will keep expanding.

The second is that the residual is itself contracting. The supply of family caregivers depends on adult children per older adult, a ratio that has fallen since the boomers had fewer children than their parents did and will fall further as post-2007 cohorts reach middle age. Meanwhile, the population needing care rises steeply from the mid-2030s. The caregiver support ratio, by which I mean the number of adult children available per person aged 80 and above, deteriorates over exactly the period when demand peaks.

Geographic dispersion compounds it. The out-migration of Section 2.5 is also a description of adult children moving several hundred miles from parents they will eventually need to care for. **The formal sector cannot absorb the overflow because of the wage and immigration constraints of Sections 3.5 and 4.2. The informal sector cannot absorb it because the families are smaller and farther away.** This is the sharpest version of the argument in this paper, and it depends on no contested projection.

The care economy has a second end, and it is usually analyzed separately from the first for no good reason. Childcare is subject to the same Baumol dynamics, the same low wages, the same immigrant labor dependence, and the same absence of a functioning private market. Its costs exceed public college tuition in most states, and Section 2.2 identified those costs as a direct fertility constraint.

The two ends meet in the same households. Roughly one in four adults caring for an adult is also raising children. The sandwich generation is not a lifestyle description but a labor supply fact: it is how care obligations at both ends of the age distribution withdraw prime-age workers from the labor force at the moment Section 4.1 says that force stops growing.

The financial consequences for caregivers compound across a lifetime and are systematically understated because they are counted as choices. A caregiver who reduces hours loses current wages, employer retirement contributions, and Social Security credits, and the benefit formula uses the highest 35 years of earnings, so years spent caring enter the calculation as zeros or near-zeros.

The caregiver is disproportionately a woman in her fifties or sixties, that is someone in her peak earning years and within a decade or two of her own retirement, and who will on average live longer and need more care herself. **The system therefore converts unpaid care today into higher public cost tomorrow, through lower payroll tax receipts, lower earned benefits, and greater eventual Medicaid eligibility.** Loop C of Figure 1 is the formal statement of this, and Section 9 asks whether caregiver credits or direct compensation would interrupt it.

8.3 The Insurance Market That Never Formed

Between the household and the state there ought to be an insurance market. There isn't one. Roughly 70 percent of Americans turning 65 will need some form of long-term services and supports (ASPE 2022), yet only about 7.5 million adults hold standalone long-term care policies, a count that excludes hybrid life and annuity products. The private market contracted sharply after insurers underpriced the risk in the 1990s and 2000s, misjudging lapse rates, interest rates, and longevity together, and the survivors have raised premiums on legacy policies repeatedly. What remains is a thin market of hybrid products attached to life insurance or annuities, purchased mainly by people affluent enough to self-insure.

The costs the market failed to cover are high and rising. The 2025 CareScout survey puts the national median at $355 per day for a private nursing home room, about $129,575 a year; $315 per day semi-private, about $114,975; and $6,200 a month for assisted living, about $74,400 a year. A home health aide at 44 hours a week runs roughly $80,000

annually (CareScout 2026). Costs are projected to rise 4 to 5 percent per year. Against a median retiree income, these are not expenses that can be absorbed. They are events that reallocate a household balance sheet.

The result is a three-tier system that maps almost exactly onto the K. The affluent self-insure, paying from assets and preserving choice about setting and quality. The poor qualify for Medicaid immediately. The middle spends down, converting a lifetime of accumulated assets into care until it becomes poor enough for Medicaid to take over, which is why Medicaid, a program designed for the poor, is the primary payer of long-term care for the middle class.

Spend-down is not a loophole. It is the system working as designed, and it means the intergenerational wealth transfer examined in Section 8.5 will be substantially smaller for middle-income families than headline projections suggest, because a significant share of middle-class estates will be consumed by care before they can be inherited.

This bears directly on the fiscal analysis of Section 3. The Medicaid long-term care line is not a welfare expenditure in the usual sense. It is the residual social insurance that appears where a private market failed and where family capacity is exhausted, and the 2025 reconciliation act reduces it just as the 85+ population begins its steepest climb. The policy is not merely cutting a program. It is removing the last backstop in a chain where the private market has already withdrawn, and the family sector is contracting.

8.4 Who Takes Care of You When Nobody Is Left

The caregiving argument in 8.3 assumes a caregiver exists. For a growing share of older Americans, that assumption fails, and the failure is the logical endpoint of sustained low fertility rather than a separate phenomenon.

Several trends converge on the same household. Childlessness has risen across successive cohorts, so more people reach old age with no adult child at all. Sibling networks are smaller, so those without children have fewer collateral relatives. Marriage rates have fallen, and divorce among older adults has risen, so fewer people reach 80 with a spouse, and the spouse is the single most common source of care. Single-person households have grown steadily as a share of the total.

And geographic dispersion, described in Section 2.5 as young workers leaving lower-branch counties, is simultaneously a description of adult children living several hundred miles from parents who will eventually need daily help. A person can have three children and still be functionally kinless at 85 if all three live across the country and work full time.

One refinement matters, because the crude version of this argument is wrong. Kinlessness is not the same as the absence of an effective support network. Someone without a spouse or children may have siblings, nieces and nephews, stepchildren, close friends, or a religious community, and the emerging research on kinless older adults finds that outcomes track the presence of a functioning network rather than the presence of any particular relationship, with the absence of a spouse and biological children not uniformly associated with worse health (Brown et al. 2026). The demographic risk is therefore not childlessness as such. It is the simultaneous thinning of every redundant layer, spouse, children, siblings, and nearby relatives, which is what sustained low fertility across successive cohorts produces, and it arrives precisely as formal care becomes scarcer.

The question this raises has no comfortable answer under any current policy. Section 3.4 established that public long-term care provision is contracting, and Section 8.2 established that the unpaid family sector absorbs the difference. For the kin-poor household, there is no family sector. Care must be purchased in the market described in Section 8.3, which most cannot afford, or provided by Medicaid after spend-down, or not provided.

The most likely outcome is earlier institutionalization, since the marginal function of a family caregiver is often precisely what allows someone to remain at home, and institutional care costs the public system substantially more than the home care it replaces. Kinlessness therefore carries fiscal consequences, because it removes the informal care that often delays institutionalization, and it grows on a schedule set by fertility and marriage decisions made decades ago.

8.5 The Wealth Transfer and Why It Widens the K

One more household-level mechanism belongs here, because it is the channel through which the divergence this paper describes reproduces itself in the next generation.

The largest intergenerational transfer of assets in American history is underway, and it is widely misunderstood. Cerulli projects $124 trillion transferred through 2048, with $105 trillion flowing to heirs and $18 trillion to charity (Cerulli 2024). The figure is contested. A 2026 Visa analysis puts boomer transfers to Generation X and millennials at roughly $36 trillion over twenty years, and the gap reflects different definitions of the transferring population, treatment of spousal transfers, and assumptions about drawdown in retirement (Visa 2026). Under the confidence classes of Section 1 this is projection-dependent and the range is wide. What is not contested is its distribution.

More than half the total volume, roughly $62 trillion by Cerulli's accounting, comes from high-net-worth and ultra-high-net-worth households, which together are about 2 percent of all households. A further $54 trillion moves horizontally between spouses before passing to the next generation, with nearly $40 trillion going to widowed women in the boomer and older cohorts. The framing that a generation is about to inherit its way into security describes what will happen to a small minority of households.

For the median household the inheritance is modest, arrives late, and is subject to the long-term care spend-down described in Section 8.3. Roughly 70 percent of Americans turning 65 will need some long-term services and supports (ASPE 2022), only about 7.5 million adults hold standalone long-term care policies, a figure that excludes hybrid life and annuity products, and median costs above $129,000 a year convert middle-class estates into care expenditure before they can be inherited. That channel is quantitatively significant and largely absent from published analysis of the wealth transfer.

The K-shaped consequence follows. An event that transfers concentrated assets to the children of asset holders, while the children of wage earners receive little and later, widens the distribution rather than narrowing it. The assets are disproportionately held in upper-branch metros and will largely stay there.

One implication for the fertility argument deserves stating. If the constraint on family formation is the cost of housing and childcare relative to young-adult income, as Section 2.2 argued, then a transfer arriving in the recipient's fifties does nothing to relieve it. Inheritance is not a fertility policy. It arrives two decades after the decisions it might have influenced.

That observation points to a distinction the wealth transfer literature largely ignores, and it may matter more than the headline totals. Inheritance is not the only channel through which family capital moves. Affluent families increasingly transfer during life rather than at death: down payments, tuition, childcare subsidy, and capital for business formation. These arrive in the recipient's twenties and thirties, which is exactly when they can change where someone lives, whether they own, and whether they have children. Families without accumulated assets cannot make them at any age.

The relevant inequality is therefore shifting from who inherits at sixty to who receives family capital at thirty, and the second is far more consequential for the demographic outcomes this paper tracks, because it operates on decisions that are still open. Two households with identical earnings and different parents will diverge in homeownership and completed fertility, and the divergence will be invisible in any income statistic. This is the K

reproducing itself across generations through a mechanism inheritance data does not capture.

8.6 Where That Leaves Us

Initially, the care economy is not mainly a labor market or a funded program; rather, it is a household activity. Unpaid family care alone has a replacement value that surpasses Medicaid spending, acting as the residual claimant that absorbs all constraints imposed elsewhere in the system. Second, that residual is contracting at the moment demand peaks, because families are smaller and more dispersed and a growing share of older adults have no available caregiver at all, so the absorption that has made the system appear to work is ending on a schedule set by decisions made decades ago.

Third, the private long-term care insurance market that might have let households self-insure never formed, so the middle spends down into Medicaid and the public program designed for the poor is the primary long-term care payer for the middle class. Section 9 turns to what could actually be done.

9. What Would Actually Help

This paper has described four reinforcing loops and a set of institutions designed for growth that are now absorbing contraction. This section asks which interventions interrupt those mechanisms, and which merely relocate their costs. It does not attempt a complete policy program, and it separates two questions public debate tends to merge: whether an intervention would work, and whether it could pass. The first is analytic and the second is not. Conflating them causes analysts to rule out effective policies on political grounds and to recommend ineffective ones because they seem achievable.

9.1 How I Am Judging These

Three tests apply to each option below.

- Does it enter a loop at a point with leverage, or does it act somewhere the cycle will absorb?
- Does it expand real capacity, or does it move a cost from a visible budget to an invisible one, typically a household?
- And how long is its lag between enactment and effect, measured against the timeline in Section 10?

The third test does most of the sorting. An intervention with a twenty-year lag is not a response to a 2032 deadline, however sound it may seem on the merits. Table 10 applies the three tests to the options below. **The pattern it reveals is that leverage and political room run inversely: the interventions that expand capacity fastest have the least political room, and the ones that are politically available mostly reallocate claims rather than create capacity.**

Table 10. Interventions assessed against the three criteria

Intervention	Loop addressed	Time to effect	Real capacity effect	Principal constraint
Immigration levels and sector channels	C, partly A and B	1 to 5 years	High; the only lever that moves labor supply within a decade	Political; integration capacity; tension with the wage strategy below
Direct care compensation	C	1 to 3 years	High if supply responds; the binding question is elasticity at these wages	Runs through Medicaid reimbursement and state budgets; sector wage pressure
Training and credentialing capacity	C and D	3 to 8 years	Medium to high, and cumulative	The supplying institutions are the ones contracting; enrollment-based funding cannot target it
Housing supply and permitting reform	A and B	5 to 20 years	High but slow; serves both ends of the age distribution	Local jurisdiction and concentrated local opposition

Care technology and assistive devices	C	Uncertain	Medium; extends independent living more than it reduces hours of presence	Task characteristics resist automation; Baumol conditions
Public employment in care	C	2 to 5 years	High; sets a floor where the dominant purchaser has suppressed one	Administration at scale; quality assurance; inflation transmission
Caregiver credits and compensation	C	1 to 3 years	Low direct capacity gain; prevents capacity loss from labor force exit	Cost; does not create caregivers where none exist
Pronatalist cash incentives	A	20+ years	Low; affects timing more than completed fertility	Lag exceeds the horizon of every deadline in Section 10
Social Security financing reform	Fiscal, not a loop	Immediate to long	None directly; reallocates claims on output	Distributional, and the phase-in window closes in 2032

Productivity and AI diffusion	Broad	Uncertain	Potentially high in aggregate; near zero in direct care	Sector mismatch; the shortage is in work AI does not touch

Source: author. Capacity ratings are the author's qualitative assessments, not measured results; each depends on recruitment, training, displacement, administration, and funding that this paper does not model. Time to effect measures the lag between enactment and material change in the relevant quantity, not the lag to full implementation. Capacity effect is assessed against the productive capacity frame of Section 3.1 rather than against budget cost.

9.2 The Zero-Sum Trap

One feature of contraction shapes every option below, and it needs stating before the options themselves. During the growth era, American states could all grow at once. Some grew faster than others, and the competition among them for people and firms was real, but it took place inside an expanding total. A state could gain without another losing.

That condition has ended. With national population growth approaching zero and turning on immigration alone after 2030, interstate competition for residents becomes arithmetic subtraction. When Tennessee attracts a retiree from Illinois, the country has not gained a retiree. When Texas recruits a manufacturer from Ohio, national output may rise or fall depending on productivity, but national headcount does not move at all. What was a race has become a division of a fixed pool.

The instruments are familiar because they are already in use, and they are expanding. States exempt retirement income from taxation to attract retirees, which works for the state and costs the nation nothing but forgone revenue. Localities offer relocation incentives and abatements to move employers across state lines. Public universities discount tuition to out-of-state students, competing for a national pool of eighteen-year-olds that Section 5.3 shows is shrinking by thirteen percent. Health systems in shortage states raise wages and offer bonuses to recruit nurses trained elsewhere, which relieves one shortage by deepening another. Each of these is rational for the jurisdiction that does it. None of them adds a single worker, student, or caregiver to the country.

This is the domestic form of a beggar-thy-neighbor policy, and the analogy is exact in the way that matters. In the trade version, a nation tries to fix its domestic problems, such as

high unemployment or a weak economy, at the expense of its trading partners by imposing tariffs and other trade restrictions.

In the domestic version, states bid against one another for a fixed population, the bids escalate, and bidding wars of this kind can dissipate fiscal capacity without expanding national headcount. The bidding is not free. Every dollar of retiree tax exemption and every relocation abatement is a dollar not spent on the training capacity, the care wages, or the housing supply that would have expanded the pool rather than redistributed it.

Two qualifications belong here, because the strong form of this claim would be wrong. Internal migration is not zero-sum in welfare terms. Moving a worker from a low-productivity location to a high-productivity one raises national output, and a substantial literature on agglomeration establishes that the gains can be large. The zero-sum claim holds for headcount and for the fiscal base, not for output. And some interstate competition is genuinely productive: it permits regulatory experimentation, and it lets households sort toward the tax and service mix they prefer.

The argument here is narrower. Competition that was largely productive under growth becomes substantially rent-seeking under contraction, because the margin on which states compete shifts from creating advantages to capturing population that already exists.

The framework explains something the rest of this section would otherwise leave as a puzzle. Section 9.9 observes that the interventions with the most leverage consistently have the least political room and treats that as a fact about politics. It is better understood as a structural consequence. Competing for share is cheap, fast, and visible: an abatement produces a ribbon-cutting inside one electoral cycle. Expanding national capacity is expensive, slow, and leaky: a state that raises care wages or funds nursing programs finds that trained workers are mobile and its neighbors benefit from the investment. Any individual governor faces that calculus and reaches the same conclusion, which is why the pattern is uniform across states of every political complexion.

The structure is the same one Section 4.4 identifies in firms. No firm captures the full return on training a junior worker, so all firms underinvest in training even though the aggregate return is high. No state captures the full return on expanding national capacity, so all states compete for existing population instead. It is a coordination failure operating at two levels at once, and in both cases the individually rational choice produces the collectively worse outcome.

One implication runs through everything that follows and is worth reinforcing here. **Immigration is the only lever assessed in this section capable of increasing the working-age population directly and materially within a few years.** Training, higher care wages, older-worker participation, technology, and housing reform can all expand effective capacity, and Table 10 says so. What none of them can do on that timescale is add people. Interstate competition for residents, by contrast, moves a fixed pool between jurisdictions without adding to it. Immigration is a different kind of policy from the others on the list, and Section 9.3 takes it first for that reason.

9.3 Immigration

Immigration has the shortest lag and the largest effect of any lever we have, and this paper has documented the point from several directions. It moves labor force growth within a single business cycle, while fertility operates on a twenty-year lag. It directly supplies the sector with the most acute shortage, since roughly a quarter of direct care workers are foreign-born. It affects trust fund arithmetic materially, as the 2026 Trustees revision demonstrated. And it determines the magnitude of the reapportionment shift analyzed in Section 6.1.

The analytic case is not an argument for any particular level, and this paper takes no position on the aggregate number. The narrower observation is that immigration policy is currently made almost entirely on other grounds, and the labor market and fiscal consequences documented here are rarely part of the debate. Sector-specific channels, of the kind several countries use to recruit care workers with defined training and portability requirements, address the shortage in Section 3.5 more precisely than aggregate levels do, and would be a partial measure even under a restrictive overall regime.

The counterargument deserves noting: expanding immigrant supply into a low-wage sector can potentially suppress wage growth in that sector, which works against the compensation strategy in 8.4. Those two interventions are in tension, and a serious program would have to sequence them, raising the wage floor before or alongside expanding the channel rather than after.

9.4 The Care Economy

The care economy is where the leverage is, because Loop C runs through it, because Section 3.5 identified it as the binding capacity constraint, and because Section 8.2 showed it functions as the residual absorbing failures elsewhere. Four families of intervention are available, and they are complements rather than alternatives.

Compensation comes first and the evidence supports it most directly. Low direct care pay contributes to the turnover documented in Section 3.5, which that produces most of those openings. Raising effective wages requires raising Medicaid reimbursement, because Medicaid sets the price for roughly half the sector, which means the intervention runs through state budgets and federal matching rates rather than through labor market policy. The 2025 reconciliation act moves in the opposite direction.

The counterargument is cost: raising the wage floor across a workforce of several million is expensive, and under a conventional fiscal frame that expense competes with the benefit payments the same budget is trying to preserve. Under the capacity frame of Section 3.1, the calculation differs, because the alternative to paying for care is not saving money but doing without care, which reappears as household cost, foregone earnings, and eventual Medicaid liability.

Credentialing and training capacity is the second, and Section 5.4 established the difficulty: the community colleges and regional institutions that supply nursing and allied health credentials are the institutions under the most financial stress, in the counties where the shortage will be worst. Loop D means this problem worsens on its own. Targeted support for training capacity in shrinking regions is unusual as education policy, because it means funding programs on the basis of projected labor demand rather than current enrollment, and enrollment-based formulas cannot do it. This is a case where the growth-era institutional design is the direct obstacle.

Technology is the third and the most uncertain. Assistive devices, remote monitoring, fall detection, medication management, and eventually some physical assistance could raise output per caregiver, and Section 3.5 conceded that a paper treating the worker-to-retiree ratio as fixed would contradict its own treatment of AI.

Direct care tasks have resisted automation for decades, and for a structural reason rather than a lack of engineering effort: bathing, feeding, transferring, and supervising require a person in the room. That is why measured productivity growth in the sector has been near zero, and why its costs rise relative to sectors where productivity does grow, the pattern Section 3.5 identifies as Baumol's cost disease. Technology is more likely to extend the period an older adult can remain at home, which avoids the most expensive form of care, than to cut the hours of human presence required once serious limitation begins. The gain is real and it is not a solution.

Public employment is the fourth. A federal job guarantee anchored in the care economy, which I have argued for elsewhere, addresses the market failure directly: it establishes a wage and benefit floor in a sector where the dominant purchaser has suppressed wages,

and it does so counter-cyclically. The demographic case for it strengthens as this paper's argument accumulates, since the sector with the largest projected labor demand in the country pays wages that keep its workers near poverty.

Because this paper adopts a functional finance frame, a reader may reasonably suspect the analysis of having been built backward from this conclusion. So the objections to a care job guarantee are stated here in full rather than confined to a footnote, and they are substantial. A public wage floor in care would raise costs for private providers competing for the same workers, with effects on private-pay pricing that are not fully predictable.

Administration would be difficult, since care work is dispersed, relational, and hard to supervise, and the federal government has limited experience directly employing at this scale.

Quality assurance in a rapidly expanded workforce is a real risk, and the history of rapid expansion in home care has included fraud.

And the inflation objection cannot be dismissed by pointing to functional finance, since the argument of Section 3.1 is precisely that real resource constraints bind: bidding up care wages in a genuinely tight labor market transmits to prices.

The discipline that keeps this honest is to run two ledgers rather than one. On the first, eliminating administrative waste releases real resources without bidding for scarce ones, so the savings are genuine and non-inflationary. On the second, raising clinical and direct care wages to attract workers into a sector with a fixed near-term supply bids for resources that are already employed, and that is where price pressure appears. Combining the ledgers in either direction misleads: counting administrative savings while ignoring wage pressure overstates the net gain, and treating all care spending as inflationary ignores the waste. The general rule is that spending pressures prices when it competes for resources already in use, and clinical labor is the clearest case of that condition in the American economy.

My assessment is that these are implementation problems rather than objections in principle, and that the alternative, continued rationing through waiting lists and unpaid family labor, has costs that are larger and merely less visible. A reader who weighs the administrative risks differently can accept every other argument in this paper and reach a different conclusion here, and the demographic analysis does not depend on this recommendation.

9.5 Housing

Housing appears in three loops and is therefore high-leverage, though slow. On the fertility side, the constraint identified in Section 2.2 is the cost of family-sized housing relative to young-adult income, which is a supply problem in the metros where jobs are. On the aging side, Section 8.1 identified a retrofit gap and a shortage of smaller units in established neighborhoods that would let older households move without leaving their communities, which would in turn release family-sized stock. These are the same policy: permitting reform that allows accessory units, smaller-format housing, and moderate density in existing neighborhoods addresses both ends simultaneously. That coincidence is unusual and worth exploiting, since interventions that serve two loops at once are rare.

The limitation is lag and jurisdiction. Housing supply responds over years, the decisions are made by local governments, and the places with the most binding supply constraints have the most entrenched opposition. Nothing here relieves the 2032 deadline. Housing is a twenty-year intervention that should have started twenty years ago, which is an argument for starting now rather than an argument for dismissing it.

9.6 Social Security and Medicare

Section 3 established that trust fund dates are legal triggers rather than measures of capacity, and that the distributional question is the substantive one. Table 3 set out the Social Security options and their incidence. Two observations belong here rather than there.

The first concerns sequencing. Because phase-in time is itself a depleting resource, the menu narrows every year toward options that take effect immediately, which means options falling on current beneficiaries. An early agreement on a modest combination of measures dominates a later agreement on a larger single measure, purely on incidence grounds and independent of anyone’s preferences about the shape of the fix.

The second concerns Medicare, where the analysis in Section 3.3 points somewhere other than benefit design. If Parts B and D drive the trajectory, and if per-beneficiary cost growth rather than enrollment drives Parts B and D, then the leverage is in pricing rather than in eligibility or cost sharing. Raising beneficiary cost sharing addresses the symptom while transferring cost to exactly the lower-branch retirees for whom premiums already consume a rising share of a Social Security check. A demographic framing of Medicare growth will produce the wrong prescription.

9.7 Four Things That Will Not Work

Four responses come up constantly and none survives the evidence assembled here. Ruling out ineffective options is how the feasible set gets defined.

Existing evidence gives little reason to expect politically plausible cash incentives, on their own, to restore fertility to replacement. The international record shows cash incentives moving birth timing more than completed fertility. South Korea has spent more than $200 billion on low-birth-rate policy since 2006 by its government's own accounting, and still records the lowest fertility among advanced economies, at 0.80 in 2025 (Stangarone 2023).

Japan roughly tripled its family-policy budget over the past decade and recorded a tenth consecutive annual decline, to 1.14 in 2025 (MHLW 2026). And Section 2.2 established that by the 2040s the cohorts that would have to produce the additional births are themselves too small for the arithmetic to work. Family policy has strong justifications on other grounds, including the wellbeing of the children who are born; reversing the demographic trajectory is not among them.

Raising the retirement age as a technical adjustment to longevity does not do what its description claims. Longevity gains have accrued disproportionately to the top of the income distribution, so a uniform increase in the claiming age is a benefit cut whose severity varies inversely with income and with the physical demands of the job. It is a legitimate option with an honest label. The label usually attached to it is not honest.

Growth optimism does not resolve the arithmetic. Section 4.1 showed that holding trend output growth at historical rates while labor force growth approaches zero requires demanding and usually unstated assumptions about productivity, hours, quality, or capital intensity. Projections combining demographic stagnation with normal growth are not wrong in principle, but they carry an assumption that should be stated and defended rather than inherited.

And artificial intelligence will not solve the care shortage. Section 4.5 set out why the surplus AI creates does not fill the shortage demography creates: the workers, the skills, the wages, and the geography all differ. Technology that automates codifiable cognitive work is close to orthogonal to a shortage in work defined by physical presence. AI may raise national output and therefore the resources available to fund care, which matters under the capacity frame. It will not put a person in the room.

9.8 Where the Positions Actually Overlap

Section 9.2 explained why jurisdictions competing for a fixed population favor instruments that capture share over instruments that expand capacity. That argument concerns the

incentives facing any single state, and it says nothing about whether a national constituency exists for the capacity-expanding measures. This subsection asks that narrower question: taking the positions currently held by identifiable currents in American economic policy debate as given, where do they overlap and where do they not? The exercise is descriptive. It takes no position on which coalition should form, which movement is correct, or what any reader should prefer.

Five currents can be identified with reasonable clarity from published platforms and policy documents, and I use each group's own self-description. The Democratic Socialists of America advance economic democracy and universal provision. The Congressional Progressive Caucus pursues reform within existing institutions, including industrial policy and expanded social services. The America First Policy Institute frames policy around border security, energy production, industrial revival, and national strength. American Compass argues for conservative economics organized around families, communities, and productive capacity. The movement associated with the 2024 Republican platform combines economic nationalism with cultural restoration. I examined these five comparatively in earlier work (Lane 2025).

Table 11. Stated positions on questions bearing on demographic policy

Question	Where stated positions converge	Where they diverge
Domestic productive capacity	All five favor rebuilding domestic production and supply-chain resilience	Left currents emphasize green transition and labor standards; right currents emphasize energy production and national security
The dignity and value of work	All five hold that wage work should support a household and that this has eroded	DSA and CPC favor union empowerment; Compass favors negotiated labor-management arrangements; AFPI and the 2024 platform favor voluntary arrangements over collective bargaining

Family formation costs	Housing and childcare costs are widely identified as burdens on young families	Left currents favor universal provision and public supply; right currents favor family-targeted benefits and deregulation of housing supply
Training and workforce pipelines	Broad support for apprenticeship and vocational pathways as alternatives to four-year credentials	Disagreement over public funding levels and over the role of organized labor in training
Care work	Recognized as economically essential and undervalued across currents	Left currents treat it as a public provision question; right currents more often treat it as a family and community responsibility
Immigration	No convergence	Positions range from expansion with legalization to substantial restriction, and this is the sharpest division among the five
Federal fiscal capacity	Limited convergence in practice; all five supported large deficit-financed programs between 2020 and 2022	Sharp divergence in stated principle, with right currents generally maintaining deficit concern that their own supported legislation did not reflect

Compiled by the author from each organization's published platform, policy handbook, or legislative agenda as of 2025; the primary documents are cited individually in Lane (2025), which develops the comparison at length. Positions are as stated in those documents and may have shifted since. Stated positions are not always the same as revealed priorities. Where the table reports convergence, it means that each organization's own materials express the position, not that the organizations would agree on an instrument for acting on it. The table makes no assessment of which positions are correct.

Three observations follow, and each cuts in more than one direction.

First, the interventions this paper identifies as highest-leverage do not map onto the partisan divide the way one might expect. Raising care wages, expanding training capacity,

and increasing housing supply each draw support from stated positions on both sides, though for different stated reasons. A reader on the left may see care wages as a labor standards question; a reader on the right may see the same measure as supporting family caregiving capacity and community stability. Whether these constitute genuine common ground or merely coincident language is not something this paper can determine, and the historical record on such convergences is not encouraging.

Second, the sharpest division falls on precisely the intervention Section 9.3 identifies as having the shortest lag and the largest effect. Immigration is the one measure capable of adding working-age people directly rather than redistributing those already here, and it is also the one on which the five currents share almost nothing. The finding is uncomfortable for anyone who wanted the analysis to point toward an easy consensus, and it should be read as a constraint on what is achievable rather than as an argument for any particular immigration level.

Third, convergence at the level of principle routinely dissolves at the level of instrument. All five currents can agree that care work is undervalued while disagreeing entirely about whether the remedy is a federal wage floor, an expanded tax credit, a family caregiver benefit, or nothing at all. Agreement on the diagnosis is cheap. The evidence from the 2021 to 2022 legislative period, when several capacity-expanding measures had bipartisan verbal support and narrow or no legislative majorities, suggests that the gap between shared diagnosis and shared instrument is wide.

What the exercise establishes is narrow but not trivial. **The feasible set for the interventions in this section is not defined by the partisan divide alone. It is defined by the divide plus the zero-sum dynamic of Section 9.2 plus the timing problem of Section 9.9. A reader who concludes from this that a durable coalition is available has read more into it than the evidence supports. A reader who concludes that nothing is possible has read less.**

9.9 What Is Doable Is Not What Works

The interventions with the most leverage are, consistently, the ones with the least political room. Immigration expansion runs against the current policy direction. Raising care wages requires increasing Medicaid spending in a period of enacted reduction. Housing supply requires local action against concentrated local opposition. Public employment in care requires a fiscal frame most participants in the debate do not share. Underneath all four sits the statutory PAYGO architecture described in Section 3.1, which

makes measures that reduce claims on output self-financing and measures that expand capacity dependent on an offset or a waiver.

Meanwhile, the politically available interventions- benefit cuts phased far into the future, cost-sharing increases, and work requirements- shift costs into the household sector, where Section 8.3 showed they are absorbed silently and reappear later as reduced revenue and higher public liability.

Section 6.5 identified the mechanism that may change this. The arrangement holds while the benefit question is not salient. In the fourth quarter of 2032 it becomes salient by operation of law, on a date fixed in advance, in the counties most dependent on the transfers and least equipped to absorb their reduction. Whether that produces a durable settlement or a deferral is not something analysis can determine. What analysis can say is that the settlement will be negotiated under worse conditions the longer it waits, and that the loops described in this paper run continuously in the meantime.

10. The Consolidated Timeline

Table 12 gathers every dated milestone in this paper into a single sequence, with the confidence classes defined in Section 1. It is the paper's only milestone table; the sections above develop the entries and this one orders them. Reading it in order makes two features visible that the sectional treatment obscures. The first is the density of the early 2030s, where fiscal deadlines, the reapportionment cycle, and a presidential election converge within roughly twenty-four months. The second is that the milestones with the highest confidence, the education entries, rest on births that have already occurred, while the entries generating the most public attention, the trust fund dates, are policy-contingent by construction and describe what happens only if Congress does nothing.

Table 12. Consolidated milestone timeline, 2025-2050

Date	Milestone	Confidence class	Section
2025	Net migration collapses to a trough; estimates disagree on whether it turned negative (see 2.4); high school graduates peak at 3.8-3.9 million; unpaid family care reaches ~$1.01 trillion in replacement value	Observed (magnitude disputed)	2.4, 5.3, 8.2

2026	Enrollment cliff begins (near-certain); general revenue crosses 45 percent of Medicare financing, tenth consecutive determination and ninth consecutive funding warning (observed); Medicaid work requirements must be implemented no later than January 1, 2027 (policy-contingent)	Observed and policy-contingent	5.3, 3.3, 3.4
2027-2032	Usable window for Social Security legislation with meaningful phase-in; the feasible menu narrows each year	Policy-contingent	3.2, 9.6
2028	Medicaid cost sharing of up to $35 per service begins for certain expansion adults	Policy-contingent	3.4
2025-2029	College-age population falls roughly 12 to 15 percent; institutional closures accelerate	Near-certain	5.3
2030	All boomers 65 or older; deaths begin to exceed births, making immigration the sole source of growth; decennial census	Near-certain and projection-dependent	2.1, 2.3, 6.1
2031	K-12 enrollment falls to ~46.9 million, 7.6 percent below the 2019 peak, with 40 states below 2022 levels	Near-certain	5.1
2031-2032	First redistricting cycle under the post-Callais standard; Medicaid provider tax cap completes its decline to 3.5 percent	Near-certain (timing); policy-contingent (substance)	6.2, 3.4

2032	New House map and electoral votes take effect; OASI depletes in the same quarter as the presidential election, with 78 percent of benefits payable	Policy-contingent	6.1, 3.2
2033	Medicare HI depletes in Q2; revenue covers 89 percent of costs, an automatic cut of 11 percent rising toward 16	Policy-contingent	3.3
2034	Combined OASDI depletion if Congress merges the funds; 83 percent payable. Participation reaches 61.1 percent, ~4.3 million fewer participants than at 2024 rates. Direct care completes a decade with 9.7 million openings	Policy-contingent and projection-dependent	3.2, 4.1, 3.5
Early 2030s	Older adults outnumber children for the first time in U.S. history	Projection-dependent	2.1
Mid-2030s	The 85+ population begins its steepest climb and Medicaid long-term care demand accelerates against a falling caregiver support ratio (near-certain); millennials and Gen Z become a majority of the eligible electorate by headcount (projection-dependent)	Near-certain and projection-dependent	3.4, 8.2, 6.3
By 2035	Households headed by someone 65 or older reach roughly 50 million	Near-certain	8.1
Later 2030s	Turnout-weighted electoral crossover, lagging the headcount crossover	Projection-dependent	6.3
2040-2043	Second census and reapportionment; regional shifts compound	Near-certain (timing)	6.1

2041	High school graduates fall 13 percent from the 2025 peak, roughly 576,000 fewer	Near-certain	5.3
2040s	Under low immigration, total population peaks and begins to decline	Policy-contingent	2.4
Mid-2040s	No single racial or ethnic group constitutes a majority of the population	Projection-dependent	2.6
2050	Medicare reaches 6.5 percent of GDP, up from 3.9 in 2025; the 65+ share reaches roughly 23 percent	Projection-dependent	3.3, 2.1

Confidence classes are defined in Section 1. Near-certain entries rest on cohorts already born. Projection-dependent entries are model outputs sensitive to disputed assumptions. Policy-contingent entries describe outcomes under current law and assume no legislative action. Dates have historically moved one to three years between annual projection vintages.

11. Conclusion

The United States is entering the first sustained demographic deceleration in its history, with natural increase ending, younger cohorts shrinking, and labor force growth approaching zero, and it is doing so with a set of institutions built during eight decades of expansion. Each of those institutions carries a growth assumption in its mechanics, and each fails in a characteristic way when the assumption stops holding. School funding formulas cut revenue faster than cost. A pay-as-you-go retirement system strains as its worker-to-beneficiary ratio falls. A fixed-size legislature converts differential growth into transfers of power. Colleges depend on a replenishing pipeline of eighteen-year-olds.

The central task through 2050 is not slowing the demographic change, which is not possible, but adapting institutions that were never designed to shrink.

The burden moves downward through a hierarchy in which each recipient has less capacity to absorb it than the institution passing it along, and that task will be attempted unevenly. On one branch, aging is a retrofit market, a portfolio question, and a set of choices among care settings. On the other, it is a hospital service line that closes, a school that consolidates, a daughter who reduces her hours, and a benefit cut that contracts the

county's aggregate income. The same national trend, refracted through unequal capacity to absorb it.

One uncertainty could change these conclusions more than any policy in this paper, and it is the least discussed. The number of Americans over 85 in 2050 is close to fixed. How many of them will need daily help is not, and Section 2.1 treats healthspan as the largest genuine uncertainty in the care demand baseline.

A meaningful compression of late-life disability would do more to relieve the constraint below than any fiscal measure available. Underneath the fiscal arithmetic sits a real constraint the ledger cannot express. Someone must be in the room. The country must fill several million projected openings in work that is physically demanding, emotionally taxing, and paid near poverty wages. They must come from a domestic cohort that is shrinking and an immigrant pipeline that has sharply narrowed and whose recovery is policy-contingent, while the family sector that has silently absorbed the excess contracts for reasons set in motion decades ago.

That is the problem the trust fund debate is standing in front of and not discussing. It is not merely a financing problem. Financing reform can reallocate claims on output, and that is not trivial. It cannot by itself produce the workers, the housing, the clinical capacity, or the productivity those claims are drawn against.

The dates are known. The cohorts are already born. Six years remain before the first hard deadline, and the arithmetic of phase-in time means that every year of delay narrows the range of humane options. The unusual feature of this crisis is not its severity but its visibility. Almost nothing here depends on a surprise. That is what makes it a policy failure if it goes unaddressed rather than a misfortune.

About the Author

Edward Lane, ASA, CFP®, is an actuary and economic policy researcher focused on federal fiscal architecture, healthcare reform, Social Security, and Modern Monetary Theory. He writes about the intersection of economic governance, political economy, and public-sector modernization.

Use of AI Tools: Large language model tools were used during the preparation of this manuscript as research and editorial assistants. They assisted with organizing material, identifying potential sources and lines of inquiry, critiquing drafts, checking internal consistency, improving exposition, and suggesting alternative wording.

The paper's conceptual framework, analytical judgments, interpretation of evidence, policy conclusions, and final editorial decisions are solely those of the author. AI-generated material was reviewed, revised, and, where appropriate, rejected or rewritten by the author before inclusion. The author reviewed cited sources against original publications or authoritative repositories wherever reasonably possible and assumes full responsibility for the accuracy of the analysis, citations, interpretations, and any remaining errors.

Generative AI tools did not exercise independent scholarly judgment and are not authors of this work.

Appendix A. Implications for Institutional Investors

This paper is written for policy readers, but several findings have direct asset implications, and the abstract names institutional investors as an intended audience. This appendix states them compactly. None is a recommendation, and all inherit the confidence classes of the underlying analysis.

The demographic drivers below are unusually well identified for long-horizon work, because the cohorts are already born. What is uncertain is timing, policy response, and the extent to which markets have priced a trend visible for years. Treat the demographic input as high confidence and the pricing implication as an open question.

A.1 A Demographic Credit Screen

The claim in Section 7.4, that these exposures share a demographic factor, is testable, and the test does not require proprietary data. The variables below are available at county and state level from the Census Bureau and from standard municipal disclosure. They are proposed here as a screen rather than applied, since building and validating one against realized credit outcomes is a separate piece of work.

Table A1. Proposed demographic credit screen

Indicator	What it measures	Why it matters for credit
Ten-year total population change	Direction and persistence of decline	Size of the revenue base and the denominator for every per capita cost
Working-age population change	Taxpayer and workforce capacity	Separates aging in place from out-migration; distinguishes a retirement destination from a failing labor market
Ratio of 65+ to working age	Service demand against contributory base	Drives Medicaid, long-term care, and age-based exemption exposure

Net domestic migration	Direction of local sorting	Leading indicator; moves before assessed values register
School-age population change	Education fixed-cost exposure	Predicts consolidation pressure and district revenue decline three to five years ahead
Own-source revenue per capita and debt service coverage	Fiscal buffer	Determines whether the jurisdiction can absorb the above or must pass it through

Source: author. All six are available from Census Bureau population estimates and standard municipal financial disclosure. The screen is proposed, not validated.

Overlaying five jurisdiction-specific factors converts the screen into a risk matrix: pension funded status, dependence on a single hospital system, exposure to a tuition-dependent college as a major employer, water and sewer debt relative to connection count, and housing market liquidity measured by transaction volume rather than price. A jurisdiction scoring poorly on the demographic screen and carrying two or more of those overlays is the profile this paper is about, and it is not a profile any single sector analysis would flag.

A.2 Asset Implications

Table A2. Demographic drivers and asset-side implications

Exposure	Demographic driver	Direction and horizon
Municipal general obligation credit	Population loss and aging tax base against fixed infrastructure and legacy pension obligations (7.1-7.3)	Widening dispersion by jurisdiction through the 2030s; the national aggregate obscures divergence and the analysis has to be done locally

Utility and water revenue bonds	Falling connection counts against systems sized for prior populations (7.2)	Rising rate pressure and coverage deterioration in shrinking service territories
Rural hospital and health system debt	Worst payer mix, oldest population, Medicaid reimbursement contraction (3.4)	Elevated distress risk from 2026 as reconciliation act provisions take effect
Higher education debt	High school graduates fall 13 percent through 2041; 100 or more institutions already at elevated closure risk (5.3)	Sustained pressure on small tuition-dependent privates and regional publics; flagships insulated
Senior housing and assisted living	85+ population nearly triples by 2060; group quarters use rises sharply after 80 (2.1, 8.1)	Demand growth near-certain; staffing cost and availability, not demand, is the binding operational constraint
Home health and direct care services	9.7 million openings 2024-2034; Medicaid sets roughly half the sector price (3.5, 4.2)	Volume growth against administratively constrained pricing and severe labor scarcity; margin depends on reimbursement policy
Residential construction and retrofit	Roughly 50 million households aged 65+ by 2035 in stock lacking basic accessibility (8.1)	Durable accessibility retrofit demand, concentrated where households hold equity to fund it
Housing in shrinking counties	Out-migration, natural decrease, school and hospital closure (2.5, 5.1)	Thin transaction volumes and impaired price discovery; illiquidity risk exceeds price risk

State fiscal divergence	Differential migration, fertility, and legacy obligations (7.5)	Widening credit dispersion across states; the relevant comparison is increasingly state versus state rather than sector versus sector
Aggregate growth assumptions	Labor force growth approaching zero (4.1)	Long-horizon return assumptions built on historical trend growth embed a productivity assumption that should be stated explicitly

Source: author. Section references indicate where each driver is developed. This table describes demographic exposure, not valuation, and takes no position on whether these trends are priced.

One cross-cutting observation deserves emphasis. Municipal credit, rural hospital debt, higher education debt, and shrinking-county housing are analyzed by different specialists using different frameworks and are conventionally treated as unrelated exposures. They share a demographic risk factor that cuts across all four. If that factor dominates, they deteriorate in the same counties, in the same decade, for the same reason. A portfolio diversified across those sectors may be considerably less diversified than it appears, and the correlation will become visible only after it matters.

This report is for informational and educational purposes only. It is not investment advice, an offer, or a solicitation to buy or sell any security. All investments involve risk, and you should consult a qualified financial professional regarding your unique circumstances.

References

AARP and National Alliance for Caregiving. 2025. Caregiving in the US 2025. Washington, DC: AARP.

ASPE (Office of the Assistant Secretary for Planning and Evaluation). 2022. What Is the Lifetime Risk of Needing and Receiving Long-Term Services and Supports? Washington, DC: U.S. Department of Health and Human Services.

Auclert, Adrien, Hannes Malmberg, Matthew Rognlie, and Ludwig Straub. 2025. "The Race Between Asset Supply and Asset Demand." NBER Working Paper 34470. Cambridge, MA: National Bureau of Economic Research.

Audoly, Richard, et al. 2026. "Do Job Postings Show Early Labor-Market Effects of AI?" Liberty Street Economics. New York: Federal Reserve Bank of New York.

Baumol, William J., and William G. Bowen. 1966. Performing Arts: The Economic Dilemma. New York: Twentieth Century Fund.

Brennan Center for Justice. 2026. How States' Seats in the U.S. House Could Change After the Next Census. New York: Brennan Center.

Brown, Susan L., I-Fen Lin, Francesca A. Marino, and Kagan A. Mellencamp. 2026. "Kinlessness, Sole Family Survivorship, and the Mental and Physical Health of U.S. Older Adults." Social Science & Medicine 393: 118996.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working paper. Stanford Digital Economy Lab.

Bureau of Labor Statistics (BLS). 2025. Employment Projections, 2024-2034. Washington, DC: U.S. Department of Labor.

CareScout. 2026. 2025 Cost of Care Survey. Richmond, VA: CareScout, a Genworth company.

Case, Anne, and Angus Deaton. 2020. Deaths of Despair and the Future of Capitalism. Princeton, NJ: Princeton University Press.

CDC. 2024. "About Healthcare Workers." *Healthcare Workers*. https://www.cdc.gov/niosh/healthcare/about/index.html.

Cerulli Associates. 2024. Cerulli Anticipates $124 Trillion in Wealth Will Transfer Through 2048. Boston: Cerulli Associates.

Cervas, Jonathan. 2026. 2030 Apportionment Projections Based on 2025 Population Estimates. Unpublished calculation, Carnegie Mellon University. https://x.com/RedistrictNet/status/2016161865550659995?s=20

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. 2016. "The Association Between Income and Life Expectancy in the United States, 2001-2014." Journal of the American Medical Association 315 (16): 1750-66.

Chidambaram, Priya, Alice Burns, Tricia Neuman, and Robin Rudowitz. 2025. "5 Key Facts About Nursing Facilities and Medicaid." San Francisco: KFF.

Committee for a Responsible Federal Budget (CRFB). 2026. Analysis of the 2026 Medicare Trustees' Report. Washington, DC: CRFB.

Congressional Budget Office. 2025. An Update to the Demographic Outlook, 2025 to 2055. Washington, DC: CBO.

Congressional Budget Office. 2026. The Demographic Outlook: 2026 to 2056. Washington, DC: CBO.

Congressional Research Service (CRS). 2010. The Statutory Pay-As-You-Go Act of 2010: Summary and Legislative History. CRS Report R41157. Washington, DC: Library of Congress.

Edelberg, Wendy, Stan Veuger, and Tara Watson. 2026. "Macroeconomic Implications of Immigration Flows in 2025 and 2026: January 2026 Update." Washington, DC: Brookings Institution, January 13.

Election Data Services. 2025. Impact of Latest Population Estimates on 2030 Reapportionment. Manassas, VA: Election Data Services.

Fikri, Kenan, Sarah Eckhardt, and Benjamin Glasner. 2024. "The Great 'Transfer'-Mation." https://eig.org/wp-content/uploads/2024/09/Great-Transfermation.pdf.

Fitch Ratings. 2025. "Policy Headwinds Weigh on U.S. Higher Education Outlook." New York: Fitch Ratings, December 4. https://www.fitchratings.com/research/us-public-finance/policy-headwinds-weigh-on-us-higher-education-outlook-04-12-2025.

Fitch Ratings. 2025. U.S. Higher Education Sector Outlook 2026. New York: Fitch Ratings.

Grawe, Nathan D. 2018. Demographics and the Demand for Higher Education. Baltimore: Johns Hopkins University Press.

Houser, Ari, Rita Choula Caldera, Brendan Flinn, and Susan Reinhard. 2026. Valuing the Invaluable 2026: Family Caregivers' Contribution Reaches $1 Trillion. Washington, DC: AARP Public Policy Institute.

Humlum, Anders, and Emilie Vestergaard. 2025. "Large Language Models, Small Labor Market Effects." NBER Working Paper 33777. Cambridge, MA: National Bureau of Economic Research.

Institute of International Education (IIE). 2025. Fall 2025 International Student Enrollment Snapshot. New York: IIE.

Joint Center for Housing Studies of Harvard University (JCHS). 2016. Projections and Implications for Housing a Growing Population: Older Households 2015-2035. Cambridge, MA: Harvard University.

Kelton, Stephanie. 2020. The Deficit Myth: Modern Monetary Theory and the Birth of the People's Economy. New York: PublicAffairs.

KFF. 2025. Allocating CBO's Estimates of Federal Medicaid Spending Reductions Across the States: Enacted Reconciliation Package. San Francisco: KFF.

KFF. 2026a. Key Facts About Medicare Spending Trends and Projections from the 2026 Medicare Trustees Report. San Francisco: KFF.

KFF. 2026b. "Who Are Direct Care Workers, and How Might Federal Policy Changes Impact the Workforce?" San Francisco: KFF.

kffpriyac. 2025. "5 Key Facts About Nursing Facilities and Medicaid | KFF." *KFF*. https://www.kff.org/medicaid/5-key-facts-about-nursing-facilities-and-medicaid/.

Kinge, Jonas Minet, Jonas Christoffer Modalsli, Simon Øverland, Håkon Kristian Gjessing, Mette Christophersen Tollanes, Ann Kristin Knudsen, Vegard Skirbekk, Bjørn Heine Strand, Siri Eldevik Håberg, and Stein Emil Vollset. 2019. "Association of Household Income With Life Expectancy and Cause-Specific Mortality in Norway, 2005-2015." Journal of the American Medical Association 321 (19): 1916-25.

Lane, Edward. 2024. “Saving Social Security.” Levy Economics Institute Working Paper 1042. Annandale-on-Hudson, NY: Levy Economics Institute of Bard College. https://www.levyinstitute.org/publications/saving-social-security/.

Lane, Edward. 2025. “How Modern Monetary Theory Can Bridge the Gap Between Left- and Right-Wing Populism.” SSRN working paper 5623732. https://ssrn.com/abstract=5623732.

Lane, Edward. 2026a. “The K-Shaped Economy: Divergent Fortunes Across American Counties.” SSRN working paper, https://papers.ssrn.com/abstract=6377838

Lane, Edward. 2026b. “The 2030 Census and American Political Economy.” SSRN working paper, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6689458

Lane, Edward. 2026c. “Medicare for All Who Want It.” https://edwardlane.substack.com/p/medicare-for-all-who-want-it-what?r=zh7v.

Lerner, Abba P. 1943. “Functional Finance and the Federal Debt.” Social Research 10 (1): 38-51.

lori@accutrain.com. 2025. “Projection: Public School Enrollment Expected to Lose Nearly 4 Million Students by 2031 - AccuTrain.” AccuTrain. August 7, 2025. https://accutrain.com/projection-public-school-enrollment-expected-to-lose-nearly-4-million-students-by-2031/.

Louisiana v. Callais, 608 U.S. ___ (2026).

Massenkoff, Maxim, and Peter McCrory. 2026. “Labor Market Impacts of AI: A New Measure and Early Evidence.” Anthropic Economic Index. Early evidence report.

Medicare Trustees. 2026. 2026 Annual Report of the Boards of Trustees of the Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds. Baltimore, MD: Centers for Medicare & Medicaid Services.

MHLW (Ministry of Health, Labour and Welfare, Japan). 2026. Vital Statistics of Japan: Preliminary Figures for 2025. Tokyo: MHLW.

MPI (Migration Policy Institute). 2026. Population Aging Is Here. Immigration Can Offer Half an Answer. Washington, DC: Migration Policy Institute.

Mudde, Cas, and Cristóbal Rovira Kaltwasser. 2017. Populism: A Very Short Introduction. Oxford: Oxford University Press.

National Center for Education Statistics (NCES). 2025a. Projections of Education Statistics to 2031. Washington, DC: U.S. Department of Education.

NCES. 2025. Common Core of Data: Public Elementary/Secondary School Universe Survey. Washington, DC: U.S. Department of Education.

NCHS (National Center for Health Statistics). 2026. "Births: Provisional Data for 2025." By Brady E. Hamilton, Michelle J. K. Osterman, and Elizabeth C. W. Gregory. Vital Statistics Rapid Release. Hyattsville, MD: NCHS.

OECD. 2025a. OECD Employment Outlook 2025: Setting the Scene: Demographic Change, Economic Growth and Intergenerational Inequalities. Paris: OECD Publishing.

OECD. 2025b. Shrinking Smartly and Sustainably: Strategies for Action. OECD Rural Studies. Paris: OECD Publishing.

Penn LDI (Leonard Davis Institute of Health Economics). 2025. Over 500 U.S. Hospitals Have Stopped Delivering Babies Since 2010. Philadelphia: University of Pennsylvania.

Pew Charitable Trusts. 2026. The Aging Population Is Changing States' Revenue, Spending, and Service Demand Outlook. Issue brief, July. Philadelphia: The Pew Charitable Trusts.

PHI. 2025. Direct Care Workers in the United States: Key Facts 2025. Bronx, NY: PHI.

Pritchett, Lant. 2024. "Immigration Is Essential and Impossible." Working paper. https://lantpritchett.org/immigration-is-essential-and-impossible/.

Public Law 111-139. 2010. Statutory Pay-As-You-Go Act of 2010.

Public Law 119-21. 2025. Budget reconciliation act of 2025.

Social Security Trustees. 2026. The 2026 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. Washington, DC: Social Security Administration.

Stangarone, Troy. 2023. "Korean Policies to Reverse the Decline in the Fertility Rate Part 1: Balancing Work and Family - Korea Economic Institute of America." *Korea Economic Institute of America*. https://keia.org/analysis/korean-policies-to-reverse-the-decline-in-the-fertility-rate-part-1-balncing-work-and-family/.

Statistics Korea. 2026. Preliminary Results of Birth and Death Statistics for 2025. Daejeon: Ministry of Data and Statistics.

Sutradhar, Uttara, Lauryn Spearing, and Sybil Derrible. 2024. “Depopulation and Associated Challenges for US Cities by 2100.” Nature Cities 1 (1): 51-61.

Tcherneva, Pavlina R. 2020. The Case for a Job Guarantee. Cambridge: Polity Press.

U.S. Census Bureau. 2023. “2023 National Population Projections.” Washington, DC: U.S. Census Bureau.

U.S. Census Bureau. 2025. Voting and Registration in the Election of November 2024. Current Population Survey Voting and Registration Supplement. Washington, DC: U.S. Census Bureau.

U.S. Census Bureau. 2026. Vintage 2025 Population Estimates. Washington, DC: U.S. Census Bureau.

Visa Business and Economic Insights. 2026. The Great Wealth Transfer Reality Check. San Francisco: Visa Inc.

Vollset, Stein Emil, et al. 2024. “Fertility, Mortality, Migration, and Population Scenarios for 204 Countries and Territories, 2022-2100: A Forecasting Analysis for the Global Burden of Disease Study 2021.” The Lancet 403 (10440): 2057-99.

Western Interstate Commission for Higher Education (WICHE). 2024. Knocking at the College Door: Projections of High School Graduates. Boulder, CO: WICHE.

Wray, L. Randall. 2015. Modern Money Theory: A Primer on Macroeconomics for Sovereign Monetary Systems. 2nd ed. London: Palgrave Macmillan.

Zallman, Leah, Steffie Woolhandler, David Himmelstein, David Bor, and Danny McCormick. 2013. “Immigrants Contributed an Estimated $115.2 Billion More to the Medicare Trust Fund Than They Took Out in 2002-09.” Health Affairs 32 (6): 1153-60.