

THE INSURANCE EXECUTIVE'S GUIDE TO DEMOGRAPHIC DATA

SEVEN PLAYS FOR CARRIERS COMPETING
ON MARGIN AND GROWTH

Data provided by [EASI Demographics Data](#)



This guide was created by Consciously in collaboration with EASI Demographics, whose demographic data has powered the insurance sector for decades.

INTRODUCTION

The pressures bearing down on insurance right now are remaking the assumptions underneath every line of business. Cyber and climate are exposing protection gaps measured in the trillions. Household structures are shifting in ways legacy products never anticipated. Population is moving across states faster than carriers can re-staff their agency networks. AI is reshaping underwriting, claims, and pricing decisions, with the leaders pulling away from the laggards faster every quarter. None of these pressures can be addressed without demographic precision, yet most insurers are still working with inputs built for a market that's moved on.

The carriers responding to all of this fall into two groups, and the gap between them is widening fast. One group is treating demographic intelligence as the strategic input it has become, feeding it into pricing, product design, AI models, and their five-year territory and growth plans. The other is still treating it as a research-team line item, pulled when a specific question demands it and otherwise ignored. The difference is showing up where it matters: in margin, in growth, and in combined ratio.

This guide is for the first group, and the carriers working to join them. It lays out the seven plays we see paying off most consistently for insurers treating demographic data as a strategic input. They aren't exhaustive, but they're the seven we see running the most often—drawn from 40 years in the demographics business.



PLAY 1:

AUDITING THE BOOK BEFORE GROWING IT

THE PLAY:

Carriers are running demographic intelligence against their existing customers (not just against the market) to find segments that are over-priced, under-priced, or about to non-renew.

The same demographic precision that finds new customers can find the segments inside the in-force book that are mispriced relative to current data: preferred segments competitors with finer targeting are picking off, under-priced segments where loss ratios are drifting, and segments where non-renewal rates are climbing before the retention dashboard catches it. Most carriers don't see which of these they have until the quarterly retention report or the next experience review surfaces the damage.

That's a missed margin most strategic planning conversations aren't even looking for. Strategic planning right now is weighted toward acquisition (new cohorts, new geographies, new channels, new product lines), and the in-force book gets reviewed through persistency, claims, and premium data... but rarely through current demographic data.

The carriers running this audit are pulling block-group-level demographic profiles against their in-force book and comparing them against current 5-year ACS and proprietary scored data. Three questions an in-force audit answers: Which segments are now operating outside the loss-cost relativities the carrier filed? Which territories are showing loss-ratio drift inconsistent with their filed territory factor? Which segments are showing non-renewal patterns that don't match the rest of the book?



The output is two-sided. Defensively, the carrier gets a ranked list of segments to re-underwrite, re-price, or proactively retain before competitors with finer targeting take them.

Offensively, it gets a read on which existing customers are now cross-selling or bundling opportunities the next renewal cycle wouldn't otherwise surface: the auto policyholder whose household just bought a home, the homeowner whose household just added an adult driver, the multi-policy customer whose household composition has shifted enough to need a different product mix entirely.

The economics make the case on their own. Industry estimates put the cost of acquiring a new insurance customer at seven to nine times the cost of retaining one. Meanwhile, J.D. Power's 2025 U.S. Insurance Shopping Study found that 57% of auto customers shopped for a new policy in the past year—the highest rate in the study's 19-year history—and roughly 30% of shoppers switched. Carriers running demographic audits against their in-force book are catching the segments most likely to be in that 30% before competitors do, and finding the cross-sell opportunities sitting in the other 70%.

Industry estimates put the cost of acquiring a new insurance customer at seven to nine times the cost of retaining one.





PLAY 2: AUDITING RATING VARIABLES FOR STALENESS

THE PLAY:

Carriers are running demographic intelligence against their existing customers (not just against the market) to find segments that are over-priced, under-priced, or about to non-renew.

Most rating plans incorporate demographic variables that were chosen years (and sometimes decades) ago. Some still earn their place. Many have decayed without anyone noticing, because the underlying demographics have shifted and the GLM has been absorbing the loss of signal into its error term ever since.

Variable selection rarely gets revisited unless the loss ratio forces the question. By the time it does, the carrier has been pricing on outdated factors for years and has lost segmentation accuracy to competitors with fresher inputs. The fix isn't complicated in principle: re-fitting the GLM with current demographic inputs on each refresh, and testing which legacy variables are still credibility-weighted, which have collapsed in predictive power, and which new variables would have entered the model if it were being built from scratch today.

The audit produces two lists. The first is what to add. The strongest candidates for new rating variables are tied to risks that didn't exist, or weren't priced, ten years ago:

- Small-business density and remote-worker concentration are cyber-exposure proxies now, since work-from-home reshaped exposure for personal cyber endorsements and small commercial alike.
- Housing-stock vintage and ADU density are now proxies for climate exposure and homeowner-

claim frequency. In wildfire and convective-storm states, building footprint per parcel has outpaced the rating territory.

- Life-stage transitions and household composition (multi-generational households, delayed household formation, adult children moving home) are reshaping auto and life risk pools that traditional age-and-marital-status factors miss.

None of these were standard rating inputs ten years ago. Several of them are emerging as inputs now.

The strongest candidates for new rating variables are tied to risks that didn't exist, or weren't priced, ten years ago.

The second list is what to take out, and it's the harder of the two. A variable that's been in the rating plan for a decade carries organizational weight beyond its predictive value: it's in the filing, the rate manual, the agent-facing quoting tools, and the actuarial memos justifying the last three rate revisions. Removing it is a regulatory and operational exercise as much as an actuarial one. The carriers doing this well are treating variable removal as a planned filing project, sequenced alongside the additions, supported with the credibility analysis that shows the variable has lost its signal, and coordinated with the states where the filing will face the most scrutiny.

The output is operational. Faster speed-to-market on rate revisions because the GLM is doing more of the segmentation work upfront and the filing is defending fewer dead variables. Tighter loss-ratio control because the factors in the rate plan are actually predicting losses. A clear inventory of what's no longer earning its place, which becomes the spine of the next filing rather than a fire drill the next time a state asks.

And over the next five years, carriers running this play keep the segments competitors with stale factors are otherwise picking off.

PLAY 3: PRICING THE HOUSEHOLD AS THE UNIT OF ANALYSIS

THE PLAY:

Carriers are coordinating auto, home, and life pricing through a shared demographic data layer.



The same household generates correlated risk signals across auto, home, and life; and the carriers reading them together are pricing better than the carriers reading them apart. Most multi-line carriers don't run that way. They run each line as a parallel function: different rating plans, different actuarial teams, different vendors, different demographic inputs. Each function optimizes its own segmentation. None of them sees the household.

The household is what the customer is. A 42-year-old with a 17-year-old new driver and a recent ADU build is one household generating correlated signals across auto frequency, homeowner exposure, life underwriting, and umbrella need. Carriers that price the household as four separate strangers leave segmentation accuracy on the table in every line. They also leave themselves vulnerable on retention, because inaccurate pricing produces premium volatility, premium volatility drives shopping, and the multi-line household is the carrier's most exposed segment when shopping turns into switching.

According to J.D. Power's 2025 U.S. Auto Insurance Study, just 51% of high-value customers—those with multi-policy bundles, long tenure, and higher premiums—say they will definitely renew with their current insurer. That's the lowest definite-renewal rate of any value segment, lower than medium-value (53%) and low-value (54%) customers. The most loyal segment in the book is now the most likely to leave... and the carriers losing them are losing them across multiple lines at once.

| The household is what the customer is.



A shared demographic data layer replaces three parallel vendor relationships. Pricing actuaries can query what a long-commute auto signal implies for home and life risk. Underwriting acceptance rules consider household activity across all three products as inputs to each one. Pricing precision now lives at the household level, not the policy level.

The output is two-sided. Segmentation accuracy improves in every line, because each line is now informed by signals from the others. And the multi-line household—the carrier’s most retention-vulnerable segment—is now also its most

retention-protected, because pricing precision and cross-sell timing both improve at the household level.

Bundled homeowners retain at roughly 95% versus 85% for non-bundlers, which is why losing them hurts most. That ten-point gap looks small until compounded: over five years, the bundled book keeps three-quarters of its customers, while the non-bundled book keeps less than half. This play’s payoff is keeping that retention advantage when premium volatility otherwise breaks it. Carriers that don’t run it lose the household in every line at once.



PLAY 4:

EXTENDING DEMOGRAPHIC INTELLIGENCE INTO CLAIMS

THE PLAY:

Carriers are extending their demographic data layer downstream into claims operations, where it shapes the bigger half of their combined ratio.

Demographic data inside most carriers is treated as an underwriting and pricing input. It rarely shows up in claims. That leaves most of its value on the table, because claims is where the bigger half of every premium dollar goes—roughly two-thirds, in the form of loss costs—and demographic intelligence sharpens decisions at multiple points in the claims process.

One caveat. Claims-side use of demographic data sits in more actively contested regulatory territory than pricing-side use, with state DOIs and the NAIC focused on disparate impact in fraud detection and adjuster decisions. The carriers running this play well are doing it through governed integration: documented model design, disparate-impact testing on a regular cadence, and human review on individual-case decisions. The applications below assume that governance layer is in place.

Five places extending the data layer pays back in claims:

1. Reserves accuracy.

A first notice of loss carries more signal when paired with demographic context. The same \$20,000 water-damage estimate maps to different total claim costs depending on contents value, displacement length, and additional living expense—all of which correlate with household composition and life stage. A multigenerational household with school-age children displaced for three weeks generates different ALE and contents costs than a single-occupant rental, even when the physical damage is identical. Reserves set without that context are systematically wrong in both directions, which distorts the loss pick on the individual claim and the development factors feeding the next reserve study.

2. Fraud signal layering.

P&C fraud costs the industry roughly \$45 billion annually and accounts for around 10% of all loss and loss-adjustment expenses, according to the Coalition Against Insurance Fraud. Most fraud models train on claim-history patterns. Carriers integrating AI and advanced analytics into fraud detection could generate potential savings of 20% to 40%, according to Deloitte. Household-level life-stage anomalies—claim patterns inconsistent with the household's own demographic profile and history—add a layer of signal that pure claim-history models miss.

Claims is where the bigger half of every premium dollar goes.





3. Claims routing.

Adjuster assignment is a routing decision that benefits from demographic context in a narrow way: matching case complexity to adjuster capability. A total loss involving an older homeowner, a multigenerational household, and significant ALE exposure benefits from an experienced adjuster more than a single-vehicle glass replacement does. This is workload management (matching the right adjuster to the right claim) which lifts both efficiency and customer satisfaction.

4. Subrogation portfolio modeling.

Demographic context improves the modeling of recovery probability across the subrogation book. Carriers running this layer are making sharper decisions about where to invest legal spend: which case types, which jurisdictions, which claim profiles. The recovery improvement comes from better portfolio-level allocation of effort: decisions made at the book level rather than the case level.

5. Customer experience differentiation.

Different cohorts want different claims experiences. For example, 87% of Gen Z and Millennials are comfortable managing the entire claims process digitally, while only 60% of Boomers and Pre-Boomers are. And the stakes of getting it wrong are high: 52% of customers who rate their digital claim experience as “poor” or “just OK” are likely to leave or not renew, versus only 4% of customers who rate it “excellent.” Routing by cohort means matching the channel to the customer (digital for some, agent-mediated for others, hybrid for many) rather than running one process for everyone.

The reason this play pays back faster than most is that the data layer is already built. Extending demographic data into claims is mostly an integration project, not a new build. Carriers doing this are tightening reserves and lifting claims satisfaction on infrastructure they’re already paying for.

PLAY 5:

TARGETING THE COHORT BEFORE DESIGNING THE PRODUCT

THE PLAY:

Carriers are identifying the cohort, mapping its demographics and distribution path, and finding the partners who reach it before building the product.

The four plays above are about extracting more value from the infrastructure most carriers already have. *This* play is about new business creation: using demographic intelligence to find cohorts the carrier's existing book doesn't serve.

Most carriers run product development in reverse. The product gets built first, often based on a market hypothesis or a competitor's move, and then marketing is asked to find buyers. The carrier ends up over-investing in actuarial work for products that don't have market fit, then either sunsets the product (sunk cost) or spends years on acquisition trying to force fit (acquisition spend on the wrong audience). Both are expensive, and the cost compounds across product launches because the sequencing problem stays invisible.

The cohorts most worth targeting today share a common shape: their actual exposure doesn't fit the household profile most personal lines products were built around. Multigenerational households accounted for 17% of all 2024 home purchases, yet most homeowner products still assume single-family-home dynamics (single primary insured, single ALE exposure, single insurable interest). Gig





| Most carriers run product development in reverse.

and independent workers now make up 10.2% of the U.S. labor force, with broader industry estimates running higher when supplemental gig work is counted, yet personal auto policies generally exclude commercial use, leaving a meaningful share of those drivers underinsured for their actual vehicle use. Caregiver, widowed, boomerang, and single-adult households are similarly identifiable in census data, mapped to state and metro markets, and underserved by current product portfolios.

The carriers acquiring these cohorts are doing the demographic work upstream of product design at the *block-group level*, not the zip-code level. Block-group profiling produces actionable maps of where caregiver households cluster, where multigenerational homes are growing fastest, where gig-worker concentration is highest. That precision determines which markets are worth entering and which states' filing strategies are worth prioritizing.

It also produces a different kind of business case. With demographic

precision, the carrier can quantify the gap as "X households in this state, Y of them currently underinsured for this exposure, with Z% concentrated in the block groups our distribution can reach." That's the difference between a national hypothesis and a state-by-state addressable market, and the difference between a product built once correctly and a product rebuilt three times during launch.

The distribution partner is what makes the acquisition cost work. Estate attorneys for widowed households. Gig platforms for independent workers. Agents fluent in multigenerational household dynamics. Financial advisors for caregiver households planning long-term care. The partner's existing reach inside the cohort replaces the years of marketing spend that product-first launches require.

Carriers running this sequence are reaching cohorts the rest of the industry can't easily find, through partners the rest of the industry isn't yet talking to, with products the rest of the industry hasn't yet built.

PLAY 6: TREATING INPUT QUALITY AS THE AI MOAT

THE PLAY:

Carriers are competing on what their AI is fed, not which model they've licensed.



Early AI strategy inside carriers tends to start as a model selection problem: which vendor, which agentic framework, and which platform to use. Those questions matter, but every carrier can buy the same models, on the same timeline, at increasingly similar prices. What *doesn't* commoditize is what gets fed into them, and the carriers winning on AI are pairing them with demographic inputs their competitors don't have access to.

Adoption is uneven and outcomes are starting to diverge. WTW's 2026 Advanced Analytics and AI Survey of P&C insurers in the U.S. and Canada found that carriers investing more in advanced analytics and AI ran combined ratios six points lower and grew premium three points faster than slower adopters from 2022 to 2024. McKinsey's analysis takes a longer view: over the past five years, AI leaders in insurance have generated 6.1 times the TSR of AI laggards, against two to three times in most other sectors. The carriers behind those numbers are running the same model architectures available to everyone else. The separation is coming from input quality.

Public-internet-trained AI confidently generates insurance recommendations from data that's years out of date and far too coarse for underwriting at the household or block-group level. Run on those inputs, even a state-of-the-art model returns answers calibrated to the patterns of the public internet, not to the household composition, life stage, and ZIP+4 expenditure signals that actually predict insurance outcomes.

Insurance-grade inputs have three characteristics that general-purpose data doesn't:

- **Current:** refreshed against the most recent reporting period rather than carrying multi-year lag
- **Granular:** precise enough to underwrite at the resolution the model is capable of (block group, ZIP+4, household composition rather than national or county)
- **Structured for decision use:** scored, ranked, and indexed in ways that let a model weight signals against each other rather than parsing raw fields

Building inputs with each of those characteristics takes years and requires continuous maintenance, which is what turns the input layer into a moat the longer it's run.

That moat is what the carriers running this playbook are already building. The demographic data infrastructure from the first five plays—the in-force audit, rating refresh, household pricing layer, claims integration, and cohort targeting—*isn't* separate from the AI investment. It *is* the AI investment.

The carriers winning on AI are pairing models with demographic inputs their competitors don't have.





PLAY 7: MOVING DEMOGRAPHIC FORECASTS INTO CAPITAL ALLOCATION

THE PLAY:

Carriers are moving demographic forecasting from the research function into capital allocation meetings, so strategic decisions run on forward models rather than trailing indicators.

Most carriers have demographic forecasting capability, but they *don't* have it in the meetings where strategic decisions get made. The five-year migration outlook shows up in the planning deck and disappears when agency territories get redrawn. Household-formation projections appear in product strategy memos and are absent when investment committees set portfolio weights. Aging cohort analysis gets presented to the board and ignored when M&A teams model acquisition targets. The pattern is consistent: forecasts get *cited*, but they don't get *used*.

The carriers treating demographic intelligence as strategic input have moved it from the research function into the C-suite. Their forward models inform strategic narratives and drive operational decisions.



| Forecasts get *cited*, but they don't get *used*.

In agency-network planning, they redesign territories against five-year migration corridors. When demographic models show population flowing toward specific ZIP codes, territory boundaries shift proactively and coverage capacity follows. They build agent density in growth corridors while competitors are still working from premium density maps.

In product investment, committees weight portfolio bets toward cohorts whose demographic share is expanding. They're building positions in the segments that will dominate the customer base in 2030. Their life insurance products serve dual-career households the competition hasn't designed for. Their auto coverage anticipates the urban-suburban commuting patterns emerging in fast-growing metros.

In geographic expansion, they test demographic trajectories before premium history. Market entry runs on household-formation forecasts and migration models, identifying markets where the customer base

will grow regardless of current competitive intensity. They enter markets while the demographics are still emerging as opportunities.

In M&A diligence, they model acquisition targets' books against forward population shifts to identify which segments will compound and which will decay. They're buying blocks that grow with favorable demographics while competitors chase acquisitions that require market-share gains to deliver returns.

Five years from now, the carriers making this change will own the fastest-growing customer segments because they built positions early. They'll have agency presence in ZIP codes where households are forming, while others are prospecting in markets where household formation already peaked.

Those running this play are positioning their entire enterprise ahead of demographic shifts the rest of the industry will be reacting to in 2030. The ones that *aren't* will be reacting.

SEVEN PLAYS, ONE CAPABILITY

These seven plays form a single capability applied across seven decisions: the in-force book, the rating plan, the multi-line household, claims operations, new product creation, AI, and capital allocation. The carriers running them *well* are running them together. They feed the same cleaned, ranked, and current demographic data into every function, so the precision in one place reinforces the precision in the others. Point solutions miss the compounding effect entirely.

The most competitive carriers have moved demographic data from research input to foundation: the layer under pricing, product, distribution, claims, AI, and forward planning that determines how well each works. Carriers that haven't made this

shift are losing ground in segments they haven't audited, against rating variables they haven't refreshed, in households they price as strangers, in claims they reserve without context, in cohorts their products were never built for, in AI applications running on stale inputs, and in capital decisions made on trailing indicators. The losses are combined ratio drift, share loss in growth segments, and a strategic position that compounds over time.

The carriers defining the next decade are treating demographic data the way the best once treated actuarial data: as the foundation everything else builds on. Five years from now, they'll own the segments, models, and positions their competitors will be chasing.

READY TO ACT?

Want the demographic data that powers this kind of analysis?

EASI Demographics has spent 40 years in the demographics business and 25 building the data tools insurance teams use to underwrite, target, and plan with precision.

[Get in touch with EASI →](#)

Want help building the intelligence?

Consciously's Ecosystem IQ™ Suite team works with companies to map audiences, markets, and ecosystems, turning demographic data into the kind of integrated intelligence carriers act on across pricing, product, distribution, claims, and capital decisions.

[Talk with Consciously →](#)

consciously



Data provided by EASI Demographics Data

