checkrtrust

The underwriting data waterfall〉

A new framework for insurance economics

Most carriers pay for pre-bind data on applicants who never convert. Here's a better order, and a per-closed-policy model, for personal lines.

Contents

The underwriting data cost problem is getting worse

For carriers writing personal lines (Auto, Home, and Life), the cost of underwriting data is a structural drag on margin. MVRs average roughly $11 per pull in Auto. Life insurers report paying $6 to $10 per criminal or driving record check, often as a standard step on every application. On underwriting margins that often run in the low single digits, those costs across millions of applicants compound into one of the largest recurring line items in the underwriting expense base. A carrier writing 500,000 policies per year runs data on far more applicants than that. In some channels, the applicant pool is 10 to 100 times larger than the bound book.

Two forces are widening the gap between data spend and policy yield. The first is shopping behavior. In 2025, J.D. Power found that 57% of auto insurance customers actively shopped for a new policy, the highest rate in the 19-year history of their study.[1] Home insurance shopping has climbed alongside it, with quarterly shopping rates that now run 6 to 7 percent of policyholders.[2] In Life, the MIB Life Index recorded 14.3% application activity growth in Q1 2026, the highest first-quarter growth on record.[3] Across all three lines, carriers are running per-transaction data costs on a quote and application pool that vastly exceeds their eventual book of business. The math compounds quickly: every applicant who never converts is data spend that produces no premium.

The second force is pricing structure, which is broken at both ends. The dominant incumbent maintains thousands of carrier contracts, each individually negotiated and confidential. Two carriers writing similar books pay materially different prices for the same data, and neither knows it. State MVR pass-through fees, meanwhile, keep rising. As DMVs raise prices to offset declining order volume, every carrier that cuts back makes the next pull more expensive for the carriers who didn't. The result is a self-reinforcing cycle that makes pre-bind data spend more expensive with every year.

Carriers are asking how to spend less per closed policy, cover more of the applicant pool, and surface signals that their current pre-screens miss. The answer requires changing both the data stack and the pricing model: specifically, moving from per-transaction pricing to a model in which data cost ties to bound business.

---

[1] "Despite Slowing Rate of Increase in Auto Insurance Pricing, Most Customers Still Shopping, J.D. Power Finds." **J.D. Power**, April 2025.

[2] "Quarterly Shopping LIST Report." **J.D. Power**, Q4 2025.

[3] "U.S. Life Insurance Application Activity Achieves Record-Breaking Growth in Q1-2026." **MIB**.

How underwriting data is used across insurance lines

Carriers writing Auto, Home, and Life ask different risk questions and use MVR and criminal data differently as a result:

	Auto	Home	Life
Primary risk question	Is this driver safe and legally authorized to operate a vehicle?	Does this household present elevated risk of theft, arson, fraud, or property damage?	Does this person's behavioral history indicate elevated mortality or claims risk?
MVR role	Primary underwriting input; ordered on most applicants	Not standard; often used in umbrella policies	Supplemental behavioral signal; used alongside medical and family history
Criminal data role	Pre-screen for fraud, identity risk, staged claims signals, theft and DUI	Primary entry point: peril-specific risk modeling (theft, arson, fraud)	Violent crime, DUI, behavioral patterns predictive of mortality and claims

The vast majority of Auto carriers now run a pre-MVR screen, sorting applicants before committing to a full MVR pull. The pre-screen layer answers the question of whether to pay for an MVR at all; the MVR itself delivers the driving history the carrier rates the policy on.

In Life, applicants disclose their own history, and the carrier verifies what was disclosed using whatever public record sources it has access to. As accelerated and automated underwriting expands across the industry (LIMRA has reported that three out of four U.S. life insurers run some form of automated or accelerated program),[4] public record data is doing more of the verification work as paramed exams and other traditional underwriting requirements give way to digital signals.

In Home, criminal signals are an emerging layer. Property data, claims history, and credit-based scores remain primary, but applicant-level risk like prior theft, arson, or fraud convictions surfaces patterns that property-only data cannot reveal. The use case is narrower than in Auto or Life, but it is real and growing.

The signals each line draws on differ, but the data infrastructure underneath is largely the same: a handful of incumbents, the same per-transaction pricing, and the same underlying gaps in coverage. This is why a coverage gap or a pricing inefficiency in that infrastructure shows up across Auto, Home, and Life insurance.

[4] "Life Insurers Look to Make the Underwriting Process Easier for Customers." LIMRA, March 2020.

The gaps in current pre-screen solutions

Pre-screening is now standard practice in Auto insurance, and growing in Home and Life. The underlying idea is sound: a cheaper, earlier signal filters applicants before the carrier commits to an expensive and extensive pull. But most carriers run only a single layer of pre-screen before the MVR, and three structural limitations in the solutions they use today leave real risk on the table.

1. Coverage is narrower than most carriers realize

Court record aggregation is only as comprehensive as the courthouse data feeding it, and some jurisdictions are actively reducing what's available online.[5] Industry research suggests that national criminal databases capture approximately 80% of relevant records, meaning one in five exists only at the county courthouse level, invisible to vendors relying on national sources alone. Regional carriers in some states describe meaningful gaps that national-only screens miss. The result is uneven pass rates that mask the gaps in aggregate. Life adds another layer: applicants self-disclose history that carriers can't always verify, and public record screening only backstops them when the records exist in the vendor's pipeline.

2. Records age in ways carriers rarely model for

Court data lags actual events by weeks or months, with 90-day delays common enough to affect underwriting outcomes. It matters most in Auto & Life: a recent DUI is exactly the kind of signal that should affect a mortality assessment, but it takes time to propagate from courthouse to aggregator to vendor to underwriter.

3. No single vendor brings the full picture together

A criminal records vendor surfaces offense history and misses identity conflicts. An MVR vendor surfaces driving history and misses fraud convictions. An identity vendor confirms a person exists but says nothing about behavioral risk. Every carrier running a single-vendor pre-screen is underwriting on a partial signal, regardless of which vendor they chose. The data exists, but the connection between sources doesn't.

Every carrier should be asking their current vendor: what percentage of our applicant pool carries risk signals your pre-screen isn't surfacing? Most current vendors can't answer it.

[5] "How Did South Carolina's Criminal Record Report Process Change?" DISA, February 2026.

Before the MVR:

A framework for reducing insurance underwriting data spend

The vendors who can answer that question deliver their data through a layered waterfall rather than a single pre-screen pass. Each layer answers a specific risk question and brings together signals that no single vendor surfaces alone. The framework only works if the vendors at each layer also have the coverage and data freshness their layer demands.

By the time the most expensive layer (typically a full MVR) triggers, it runs only against applicants who have passed every cheaper screen above it. Spend concentrates where it produces underwriting value.

The same waterfall applies across Auto, Home, and Life, with each line drawing on the layers most relevant to its risk model.

Pre-bind underwriting data layers

LAYER	DATA & QUESTION	APPLIES TO	COST
1	**Driver's License Status Check** *Is this license valid and active?*	Auto, Life	**Fraction** OF MVR COST
2	**Criminal & Public Records** *Does this person carry behavioral risk signals relevant to this policy type?*	Auto, Life, Home	**25–50%** OF MVR COST
3	**Recycled MVR History** *Is recent driving history already on file?*	Auto, Life	**<$1** PER CHECK
4	**Full MVR** *What is the complete current driving history?*	Auto, Life	**$2–$25** VARIES BY STATE

Layer 2 signals by line: Auto looks for fraud, theft, DUI, and staged-claim history. Life looks for DUI, violent crime, and other behavioral mortality risk. Home looks for theft, arson, and fraud. Layer 2 includes sex offender registry, DOC, warrants, and traffic records.

Auto Insurance

An applicant submits for auto coverage. Before the carrier orders an MVR, the criminal records layer surfaces a prior fraud conviction. The carrier declines—the most expensive check was never run.

Life Insurance

The criminal records layer surfaces a recent DUI before the carrier orders the MVR. The carrier routes to manual review and adjusts the premium tier rather than issuing a policy mispriced for the actual behavioral risk.

Home Insurance

A homeowner applicant's prior arson conviction surfaces in the criminal records layer. Property data, credit, and claims history would not have caught it. The layered approach surfaces it before the policy binds.

Three things make a waterfall like this work in practice: each layer must cost materially less than the next; coverage must be consistent across states and counties; and the carrier must control the rules that decide what each signal means for the underwriting decision. When those three conditions are met, the waterfall reorganizes underwriting around the data that's actually predictive of risk.

The pricing model that changes the math

The framework ensures the most expensive checks only run on applicants who need them. But it doesn't change the underlying economics of per-transaction pricing: every applicant who enters the waterfall (even one screened out at Layer 2 at a fraction of the MVR cost) is still a billable transaction. The framework lowers what each applicant costs, but total spend still rises with every applicant processed.

And in personal lines, the number of quotes is far larger than the bound book. The 57% Auto shopping rate J.D. Power reported for 2025 means carriers are running per-transaction data costs on a funnel in which most applicants never convert. Home and Life produce the same structural waste at different magnitudes.

Per-closed-policy pricing aligns data cost with policy revenue: data is paid for only when an applicant binds. Non-converting applicants cost nothing, and top-of-funnel expansion becomes effectively free from a data-spend standpoint, which lets carriers grow acquisition without the data spend dragging on margin. For Life insurance carriers, for example, per-closed-policy pricing lets public record checks run earlier in the underwriting process, surfacing risk before the carrier invests in downstream verification steps.

The economic logic is the same one carriers already apply to reinsurance and distribution: variable cost should track revenue earned, not volume processed.

Data hasn't worked this way in the U.S. market because the dominant incumbent built its business on per-transaction pricing and hasn't adopted the per-closed-policy model. The market is starting to shift around them. The magnitude of the waste is easier to see at the transaction level. Take a carrier writing 500,000 policies per year at a 3% bind rate. At $11 per MVR pull, that's roughly 16.7 million data queries to close the book. Of those, 16.2 million queries (about $178 million in annual spend) are spent on applicants who never convert. Only $5.5 million of the data spend produces a bound policy. The ratio compounds as shopping rates rise. Every point of growth in application volume adds more waste before it adds more premium.

A substantial portion of the criminal records data personal-lines carriers already use in their underwriting workflows today comes from Checkr Trust. Carriers can now access that data directly, on a per-closed-policy basis, through a single API. The waterfall structure also lets carriers stop short of a full MVR pull on most applicants, paying only for the cheaper layers when the data warrants it. The criminal coverage runs about 20% broader than incumbent providers, addressing the one-in-five criminal records that national-only databases miss. The data is also fresher: Checkr Trust refreshes its source data continuously, while reseller data typically lags by 30 days or more.

Checkr Trust's People Data Graph also resolves records that single-vendor pre-screens miss because the data alone can't be matched to a specific applicant. A court record listing "John Smith, born 1986" with a partially redacted date of birth would be too ambiguous for most vendors to use. The Graph cross-references other identifiers (address history, license records, name aliases, and PII like phone numbers and SSNs) to match the record to the right person, surfacing risk signals that would otherwise sit unused.

The data is broader, fresher, and better connected. The pricing model removes the constraint that ties data spend to funnel volume. And the gap between what carriers spend today and what they would spend under a per-closed-policy model is, for most books of business, material.

The decision

The structural inefficiencies in pre-bind underwriting data are no longer something carriers have to live with. A different infrastructure exists today: authoritative data, direct access through a single API, and pricing built around how carriers earn revenue rather than how legacy vendors prefer to bill.

A 20-minute conversation can show how a layered, per-closed-policy waterfall would change your data spend across your current book. See what your per-closed-policy data spend would look like.

BOOK A 20-MINUTE ANALYSIS 〉

checkrtrust

checkrtrust.com