Wakely
An HMA Company

Wakely.com

WHITE PAPER

Rethinking Downside Risk: The Role of Stop Loss in ACO Contracts

ACO Risk Mitigation Strategies: Part 1

Brad Heywood, ASA, MAAA

Josh Gottesman

JUNE 2026

INTRODUCTION

This paper is the first in a Wakely series for accountable care organizations (ACOs), ACO Risk Mitigation Strategies. Wakely collaborated with Josh Gottesman from Brown & Brown (B&B) to write these papers. Each installment examines one tactic an ACO can use to take downside risk while limiting its exposure to a repayable settlement with the Centers for Medicare & Medicaid Services (CMS). The audience is ACO leadership weighing two-sided risk in the Medicare Shared Savings Program (MSSP) or entry into the LEAD Model, a new total cost of care ACO model from the CMS Innovation Center (CMMI). This first paper focuses on stop loss.

The underlying problem is familiar: downside risk has become the price of admission for upside risk. MSSP has placed most ACOs onto a glide path toward two-sided risk, while LEAD requires two-sided risk beginning in the first performance year, with no minimum loss rate and first-dollar accountability for total Medicare Part A and Part B spending. Taking real risk is how an ACO earns a meaningful share of savings. The discomfort is the tail. A single difficult year, or a handful of catastrophic claimants, can become a payment back to CMS that the balance sheet was not built to absorb.

Many organizations resolve discomfort by handing the downside to a value-based care enabler or aggregator that absorbs the risk in exchange for a material share of the savings. That trade can cost an ACO as much as half of its upside (if not more) in every year that it succeeds. For organizations with scale in lives and operations that are hesitant to take on material downside risk, this paper makes the case for a different path: pairing aggregate and member-level stop-loss reinsurance with disciplined actuarial projection so that an ACO can cap its downside, protect the financial backstop CMS requires, and keep the extra upside that taking downside risk provides. We examine how each form of stop loss works, where it applies across MSSP and LEAD, and why the pricing conversation should involve both the reinsurance broker and an actuary.

GIVING THE DOWNSIDE TO AN ENABLER MEANS GIVING UP THE UPSIDE

The most common way providers manage ACO downside risk today is by avoiding it all together. They join an enabler or aggregator that takes on the risk-bearing entity role, fronts the capital and repayment obligation, and supplies the analytics, care-management infrastructure, and point of care technology. For a group new to risk or short on capital, the appeal is obvious. The enabler stands between the practice and CMS.

The cost shows up in the savings split. Reported terms vary, but enablement arrangements commonly retain up to half of an ACO's shared savings—sometimes more—in return for carrying the downside and the overhead. In a strong year, that is the single largest factor standing between providers who generated the savings and the dollars they keep. The structure offers little middle ground between success and failure. The enabler holds the risk because it holds the contract, and providers have limited ability to keep upside as they build their own capability.

Stop loss offers a way to buy down the same tail risk while keeping the contract and the savings in the ACO's hands. Rather than transferring the entire risk position to a third party, the ACO keeps its CMS relationship and purchases protection against only the outcomes it cannot absorb: an aggregate loss beyond a level it chooses, or the cost of an individual catastrophic beneficiary. The cost is a defined annual premium rather than a standing share of the upside, and remaining savings stay with the provider organization.

The Cost of the Easy Path

Handing the downside to an enabler can cost as much as half of an ACO's shared savings in every year it succeeds. Stop loss caps the same tail risk for a defined premium and leaves the upside with the ACO.

STOP LOSS, IN TWO FORMS

Stop loss for an ACO comes in two broad forms, and the distinction matters for both MSSP and LEAD.

- **Aggregate stop loss** protects the organization's overall result. It pays when the ACO's total losses against its benchmark exceed an agreed attachment point. It is available to ACOs in both MSSP and LEAD and is the primary tool for protecting the balance sheet and the required financial backstop.
- **Specific or member-level stop loss** protects against individual high-cost beneficiaries. It pays when a single beneficiary's claims exceed a deductible the ACO selects. As discussed below, its structure fits LEAD; it does not map onto MSSP's purely aggregate settlement.

Before continuing, a brief clarification is needed: both LEAD and ACO REACH already include an optional, CMS-administered stop-loss arrangement. This is a residual-based reinsurance feature in which CMS sets attachment points using a national reference population and administers the arrangement on a budget-neutral basis, such that model-wide payouts equal model-wide charges through a uniform multiplier. MSSP has no comparable feature. The commercial stop loss discussed here is different. The ACO chooses its own deductibles and attachment points, recoveries are not constrained to CMS parameters, and a specific policy can pay during the year rather than only at settlement.

AGGREGATE STOP LOSS: PROTECTING THE BALANCE SHEET

Aggregate stop loss is the foundation. It converts an open-ended downside into a known, bounded cost. The points below—developed with our co-author at B&B, Josh Gottesman—describe how ACOs and provider groups are putting it to work:

- **Pooling risk across contracts.** This is one of the least understood yet most useful options available. An organization that holds several risk contracts (e.g., multiple ACOs or a mix of Medicare and Medicare Advantage lines) can blend them into a single combined benchmark under one policy instead of insuring each arrangement on its own. Losses in one contract net against results in the others, so the policy protects the consolidated balance sheet for the year and costs materially less than separate coverage for each contract. For a board or a private equity owner, it turns a set of open-ended positions into one statement: the organization's maximum loss for the year is a known number.

- **Underwriting below the ACO level.** Aggregate coverage does not have to be written at the whole-ACO level. It can be underwritten for subgroups within the ACO, down to the underlying provider groups. This matters most in LEAD, which requires ACOs to assign at least 1% of total downside risk, on average, across their participant providers. Provider-level protection becomes a structural need, not just a preference.
- **Subgroup and NPI-level structuring.** The most granular version writes coverage at the individual NPI level for a defined subgroup of providers. This synthetically recreates a flexibility offered by ACO REACH that MSSP and LEAD do not provide. REACH scoped claims-based alignment to the tax identification number (TIN) and the individual NPI, so for professional claims, a service counted toward alignment only when both the billing TIN and the rendering NPI appeared on the ACO's participant list. A REACH ACO could have included some clinicians in a TIN and excluded others. MSSP and LEAD remove that control. Both attribute beneficiaries at the TIN level: MSSP on primary care services billed under the ACO participant TIN, and LEAD to the participant TIN that furnishes the plurality of primary care services. Therefore, a multi-specialty group that bills under one TIN cannot carve its specialists out of the attributed population the way it could in REACH. For a system that consolidates primary care and specialty care under a single TIN, that is a genuine disadvantage. In REACH, it could list only its primary care NPIs and perform well; in MSSP and LEAD—models built around primary care—the specialist activity comes along with the TIN and can potentially deteriorate results. NPI-level stop loss is a route back to that granularity. The primary care economics run as they are, while coverage is placed specifically on the specialist NPIs that create unfavorable exposure. Two conditions apply. First, the subgroup needs enough scale to underwrite credibly, often in the range of 5,000 attributed lives. Second, because CMS reports results at the ACO level rather than by NPI, the structure requires an agreed upon methodology to allocate the ACO-level result to the covered providers. Carriers already underwrite at the NPI level and actuarial work to build that allocation, drawing on detailed Medicare claims data, is what makes the structure credible and bindable.

The Specialist Carve-Out, Put-Back

REACH lets a multi-specialty group keep its specialists out of attribution. MSSP and LEAD do not. NPI-level stop loss restores the effect: let the primary care economics run, and place cover on the specialist NPIs that drive the downside.

- **Coordinating coverage with track selection.** The right amount of aggregate coverage depends on the risk track. An MSSP ACO in BASIC Level E carries a 30% loss-sharing rate, with exposure capped at 8% of participant revenue (not to exceed 4% of benchmark); an ENHANCED track ACO can owe a loss-sharing rate between 40% and 75%, with a loss recoupment limit of 15% of benchmark. A LEAD ACO in the Professional Option is liable for up to 50% of losses against its benchmark. A Global Option ACO for up to 100%, both on a first-dollar basis. The attachment point and the track election should be set together, not in sequence.
- **Attaching at or below 100% to guarantee savings.** Coverage can attach at 100% of the benchmark, or even below it—but not above the benchmark. An ACO with a 100% attachment point would guarantee no losses beyond the cost of the policy. In addition, an ACO can set an attachment point at, say, 98% of its benchmark, which transfers everything below that line to the carrier and effectively guarantees a savings result of about 2%. An ACO with a strong multi-year savings record can sometimes secure coverage at the benchmark itself, a 100% attachment point that removes first-dollar loss entirely because the carrier sees a loss as unlikely. Whether either trade is worthwhile depends on the ACO's projected distribution of outcomes, the subject of a later section.
- **Supporting the financial guarantee and repayment mechanism.** Both programs require ACOs in downside risk to post a financial backstop: a repayment mechanism in MSSP (funds in escrow, a letter of credit, or a surety bond) for ACOs in BASIC Levels C through E or the ENHANCED track, and a financial guarantee in LEAD sized at 2.0–4.0% of the prior year's Part A and B expenditures, depending on risk option and capitation election. A stop-loss policy sitting behind that obligation reduces the risk the surety or lender is underwriting. Naming the surety as loss payee on the policy can lower—and in some cases eliminate—collateral the ACO would otherwise have to tie up.

Downside Exposure and Required Backstop by Program and Track		
Program and Track	**Loss Exposure (First Dollar Once Threshold Met)**	**Required Financial Backstop**
MSSP, BASIC Level C	30% loss-sharing rate; capped at 2% of participant revenue (max 1% of benchmark)	Repayment mechanism (escrow, letter of credit, or surety bond)
MSSP, BASIC Level D	30% loss-sharing rate; capped at 4% of participant revenue (max 2% of benchmark)	Repayment mechanism (escrow, letter of credit, or surety bond)
MSSP, BASIC Level E	30% loss-sharing rate; capped at 8% of participant revenue (max 4% of benchmark)	Repayment mechanism (escrow, letter of credit, or surety bond)
MSSP, ENHANCED Track	40–75% loss-sharing rate; capped at 15% of benchmark	Repayment mechanism (escrow, letter of credit, or surety bond)
LEAD, Professional Option	Up to 50% of losses against benchmark	Financial guarantee: 2.0% of prior Part A & B (3.5% with Enhanced Primary Care Capitation)
LEAD, Global Option	Up to 100% of losses against benchmark	Financial guarantee: 2.5% of prior Part A & B (4.0% with Enhanced PCC or Total Care Capitation)

SPECIFIC STOP LOSS MEMBER-LEVEL PROTECTION FOR LEAD ACOS

Specific stop loss works at the level of the individual beneficiary. The ACO selects a deductible, and the policy pays a share of an individual's claims above it. It fits LEAD because LEAD holds the ACO accountable for total cost of care on a per-beneficiary basis, with predicted spending built from each beneficiary's risk score and the ACO's benchmark and claims data flowing throughout the year, making member-level triggers administrable. It adds little in MSSP, for a specific reason. MSSP already truncates each assigned beneficiary's annualized expenditures at the national 99th percentile, in both the benchmark and the performance-year calculations, precisely to limit the effect of catastrophically large claims. Effectively, that truncation is a built-

in, specific stop loss that applies to every ACO in the program, so a layered commercial-specific policy would be largely redundant.

- **Predictable recoveries.** Because the trigger is a defined dollar deductible on an identifiable beneficiary, recoveries are predictable rather than contingent on a year-end aggregate result.
- **ACO specific performance.** Your recoveries are not dependent on program-wide experience, but on your population.
- **Cash during the year, not only at settlement.** A specific policy can pay on claims submitted throughout the year. Even with the optional provisional settlement, LEAD financial settlement lands after the performance year ends, with member-level recoveries putting cash in hand well before that.
- **Smoother cash flow against capitation.** For ACOs taking capitation, in-year recoveries on catastrophic members help match the timing of cash in and cash out, easing the working capital strain a few large claimants can create.
- **Deductibles the ACO controls.** Unlike the CMS residual stop loss, where attachment points are set centrally, a commercial-specific policy lets the ACO choose its own deductibles and tune the level of protection to its risk appetite and capital position.

Aggregate and Specific Stop Loss Compared

	Aggregate Stop Loss	Specific (Member-Level) Stop Loss
What It Protects	The ACO's overall result against its benchmark	The ACO against individual catastrophic beneficiaries
Payout Trigger	Aggregate losses beyond a chosen attachment point	A single beneficiary's claims above a chosen deductible
Where It Applies	MSSP and LEAD	LEAD; redundant in MSSP given truncation
In-Year Cash Flow	No; follows the performance year	Yes; on claims submitted during the year
Who Sets the Threshold	Negotiated attachment point	ACO selects its own deductible
Subgroup/NPI Structuring	Yes, given scale and an agreed allocation methodology	Beneficiary-level by design
Primary Purpose	Protect solvency and the financial backstop	Predictable recoveries; bring cash forward

WHY THE PRICING CONVERSATION STARTS WITH AN ACTUARY

The most important use of an actuarial projection is not to negotiate a lower rate, but to determine whether the ACO should purchase coverage at all and, if so, which coverage. A projection of the ACO's own results—including expected benchmark, anticipated total cost of care, probability and size of a repayment, and the shape of the high-cost claimant tail—turns an abstract fear into a decision. If the work shows a meaningful chance of a shared-loss position and a policy that removes exposure is available at a defined premium, the question is no longer speculative. An ACO that has measured its downside can keep that view internal, using it to decide how much risk to retain.

Where the projection helps in the market depends on who is underwriting. The carriers who price aggregate stop loss are mostly experienced actuaries who know ACO risk, so an independent projection is less so a lever on their rate and more so a way to test it and support a particular structuring choice. Its larger effect is on the financial backstop, where the sureties and lenders behind a repayment mechanism or financial guarantee are generally less specialized in ACO economics and a credible validation does more to support the underwriting and reduce or release collateral. This work allows the two to be optimized together: an actuary can define the provider panel an ACO should put forward, and the broker can price the existing panel against the panel under consideration.

This is why the evaluation belongs to two disciplines working together. The actuary quantifies the downside, evaluated key structural choices—including how much coverage, attachment points, deductibles, and allocation methods—and supports the surety side. The broker takes that quantified picture to the market, sources carriers, and prices the policy. Wakely performs the first role and our co-author at B&B performs the second. An ACO that engages both vendors is in a far stronger position than one that hands its entire risk position and a standing share of its savings to an enabler by default.

Start with the Number, Not the Quote

Before an ACO asks what stop loss costs, it should know its own odds: what is the chance of a repayment and how large could it be? That figure determines whether to buy coverage and what coverage to buy. Pricing comes after.

WHAT TO DO NEXT

For an ACO weighing two-sided risk in MSSP or entry into LEAD, timing is critical. MSSP track elections and participant lists for the coming performance year—and the risk option, financial guarantee, and stop-loss decisions facing LEAD entrants—fall on near-term deadlines, so the analysis behind them is useful now. The sequence itself is straightforward: quantify the downside before deciding how to cover it. Model the distribution of results and the high-cost claimant tail. Evaluate aggregate stop loss to protect the balance sheet and the financial backstop and—for LEAD ACOs—member-level stop loss to make recoveries predictable and bring cash forward. Coordinate the structure with the risk track and the required guarantee. The result is holding real risk while keeping the savings that come with it and not surrendering the upside to a third party.

This paper is the first part of a broader series. Later installments, along with a webinar, will review capital management and the interplay between downside risk and the financial guarantee, the coordination of stop loss with track selection, and the distribution of downside risk to participant providers. Wakely and B&B can also provide a quick, directional read on what stop loss would cost for a given ACO's situation as a starting point. To discuss an actuarial projection of your ACO's downside and a stop-loss strategy built on it, contact **Brad Heywood, ASA, MAAA** (brad.heywood@wakely.com) or **Josh Gottesman** (josh.gottesman@bbrown.com).

ABOUT THE AUTHORS

Brad Heywood, ASA, MAAA
Senior Consulting Actuary II
brad.heywood@wakely.com
Brad is an Associate of the Society of Actuaries and a Member of the American Academy of Actuaries. He advises ACOs and provider organizations on CMS alternative payment models, including MSSP, ACO REACH, and LEAD.

Josh Gottesman
Producer, Brown & Brown
josh.gottesman@bbrown.com
Josh structures and places stop-loss reinsurance and surety bonds for risk-bearing healthcare organizations across the country.

ABOUT THIS PAPER

This paper is intended for general educational purposes for healthcare executives and provider organizations considering risk-based participation in CMS programs. It describes program features as set out in the MSSP technical specifications and application materials, the ACO REACH financial operating guides, and the LEAD Model Request for Applications available as of June 2026. CMS guidance may change, and LEAD financial methodology papers expected later in 2026 may modify several of the parameters described here. Nothing in this paper is an actuarial opinion on, or a recommendation specific to, any particular organization, and it is not legal, tax, or accounting advice. Readers evaluating stop loss should obtain an analysis specific to their own population, risk position, and contracts.

ABOUT WAKELY

Founded in 1999, Wakely Consulting Group, an HMA Company, is well known for its top-tier healthcare actuarial consulting services. With nine locations nationwide, Wakely boasts deep expertise in Medicare Advantage, Medicaid managed care, risk adjustment and rate setting, market analyses, forecasting, and strategy development. The firm's actuaries bring extensive experience across all sectors of the healthcare industry, collaborating with payers, providers, and government agencies.