The BALANCE Model Pause, the GLP-1 Bridge Extension Thru Dec 2027 & What the 80% Part D Participation Threshold Miss Signals About Medicare’s First Real Attempt to Negotiate Anti-Obesity Drug Coverage

APR 23, 2026

Abstract

- CMS paused the Part D leg of BALANCE for CY2027 via an HPMS memo on April 21, 2026, one day after the Part D parent org application deadline of April 20

- The stated reason: insufficient critical mass to clear the 80 percent NAMBA-weighted participation threshold baked into Section 2.3.1 of the March 2026 RFA

- GLP-1 Bridge (Section 402 demo) gets extended through 12/31/2027, originally scoped to run only 7/1/26 through 12/31/26

- Beneficiary price point on the Bridge stays at $50/month, operating outside the Part D benefit and payment flow

- Medicaid leg of BALANCE proceeds on track, state applications open through 7/31/26, participation start dates anywhere between 5/1/26 and 1/1/27

- Model drug list for CY2027 covers Zepbound, Mounjaro, orforglipron (pending FDA), Ozempic, Rybelsus, Wegovy, with a $245/month net price anchor on at least Zepbound KwikPen

- Cost sharing caps under BALANCE that plans never had to live with: $50/month EA and EGWP, $125/month AE/BA, $245 plus dispensing in the deductible phase, $0 catastrophic

- Narrowed risk corridor incentive (2.5 percent instead of 5 percent first threshold) apparently wasn’t enough to pull the industry across the 80 percent line

- The 340B rebate haircut of up to 5 percent, the WAC-based gross cost treatment, and the FAD PDE field all survive for future reconsideration

- Why this matters for investors, PBMs, manufacturers, RCM vendors, and anyone building infrastructure assuming Medicare anti-obesity coverage was a 2027 event

Table of Contents

What actually happened on April 21

The memo hit HPMS on April 21, 2026, which was exactly one day after the application deadline for Part D parent organizations, and anyone who has ever worked a CMS model timeline saw that sequencing and immediately understood what it meant. CMS didn't need a week to count the applications. The agency already knew by the night of April 20 whether the 80 percent NAMBA-weighted enrollment threshold had been cleared, and the speed of the response tells you the answer wasn't close. When a threshold is missed by a hair, you usually get a few days of internal debate, maybe a week of back and forth with the biggest sponsors to see if anyone is willing to revise upward. A next-day memo means the miss was material.

The text itself is measured and bureaucratic, but the substance is that the Medicare Part D leg of BALANCE is paused for CY2027 pending further evaluation and data collection. The GLP-1 Bridge, which was originally designed as a six-month on-ramp running from July 1, 2026 through December 31, 2026, gets extended a full year beyond its original sunset to December 31, 2027. Part D sponsors are explicitly

instructed not to indicate participation in the BALANCE Model within HPMS or the Bid Pricing Tool for CY 2027, which is the bureaucratic equivalent of closing the door and turning off the lights. The Medicaid side of the model continues on schedule, with state applications open through July 31, 2026, and states able to select start dates between May 1, 2026 and January 1, 2027.

The April 6, 2026 CY2027 Rate Announcement had already flagged that stakeholders were nervous. Commenters cited in the final announcement had specifically asked for the Bridge to run at least a full calendar year before BALANCE kicked in, and CMS punted that concern at the time by saying Innovation Center design was outside the scope of the rate announcement. Two weeks later the agency effectively granted the extension anyway, which suggests the internal read by early April was already that the threshold was wobbly and the fallback plan was being drafted in parallel to the application solicitation. That is not an indictment of anyone at CMMI. That is just how mature agencies run contingencies. But it does mean the pause was not a surprise to the people inside the building.

The 80 percent math and why it was always going to be hard

Section 2.3.1 of the March 2026 RFA lays out the threshold mechanics with uncommon clarity for a CMS document. The numerator is beneficiaries enrolled in Part D plans that applied to participate in BALANCE, restricted to plans included in the National Average Monthly Bid Amount calculation, so all eligible plan types plus Defined Standard plans, but excluding SNPs and EGWPs. The denominator is all Part D plans included in the NAMBA calculation, projected from February 2026 enrollment to CY2027. If that quotient lands below 80 percent, no Medicare launch.

Eighty percent is a stiff bar. The Part D market is not that concentrated if you define concentration loosely, but if you define it the way CMS was defining it for NAMBA purposes, you really only need four or five parent organizations to move in lockstep to get anywhere close. The usual suspects in terms of aggregate Part D enrollment are a familiar list: Humana, UnitedHealth, CVS Aetna, Centene, Elevance, Cigna, Kaiser in the regional plans, and a long tail. To get to 80 percent of NAMBA-weighted enrollment, the largest handful of these parents basically all had to say yes, with conviction, across both their MA-PD and PDP books.

The problem is that saying yes here was not a trivial operational or actuarial exercise. A Part D parent that opted in was committing to cover an entirely new drug class at a $50 or $125 copay, in the middle of a bid cycle where the CY2027 standalone PDP market was already wobbling from Part D redesign fallout. The basic premium

stabilization demonstration for PDPs, the $2000 out of pocket cap, the expanded federal reinsurance mechanics, the sponsor exposure in the initial coverage phase, all of that was already reshaping how plans thought about formulary risk. Dropping a GLP-1 uncapped utilization bomb into a benefit that had been redesigned to shift more risk onto plans in the first dollar range was, politely, asking a lot. Less politely, it was asking plans to volunteer to eat adverse selection in a year where they were already being asked to eat adverse selection.

CMS tried to defuse the adverse selection concern in two ways. First, the 80 percent threshold itself was supposed to mean that if the market participated, no plan got stuck being the only one holding the bag. Second, the narrowed risk corridor option (2.5 percent instead of 5 percent on the first threshold) was supposed to give plans an additional hedge. Both of those only work if the critical mass actually shows up. Once one or two of the big five decide the math doesn't pencil, the threshold becomes unreachable, and everyone else has an even stronger reason to sit out because now the adverse selection risk is concentrated rather than diluted. Classic coordination failure, the kind that voluntary CMMI models have historically struggled to escape.

What the RFA was actually asking plans to swallow

Reading the March 2026 RFA carefully, it becomes clear why a lot of plan actuaries spent the spring doing some very uncomfortable modeling work. Cost sharing was capped at $50 per 28- or 30-day supply for Enhanced Alternative plans and EGWPs during the initial coverage phase, and $125 per supply for AE and BA plans. During the deductible phase, patient exposure was capped at $245 plus a dispensing fee for a 28- or 30-day fill. The catastrophic phase sat at standard rules, which now means zero patient cost share post the Part D redesign. Participating plans had to apply these caps to every model drug from every participating manufacturer, uniformly, with no tiering discrimination across model drugs. Same tier, same cost share, no step therapy more burdensome than FDA labeling, and PA criteria no more burdensome than the specific criteria CMS defined in Section 2.2.5.

The PA criteria are worth lingering on because they are notable for what they permit rather than what they restrict. Provider attestation that the patient has type 2 diabetes, or noncirrhotic MASH with F2 to F3 fibrosis, or OSA, or is on lifestyle modification plus one of three BMI-plus-comorbidity ladders. The BMI ladder starts at 35 for unrestricted access, drops to 30 with a qualifying comorbidity (HFpEF, uncontrolled hypertension, CKD stage 3a+, moderate to severe OSA, or MASH with F2 to F3 fibrosis), and drops further to 27 with pre-diabetes, prior MI, prior stroke, or

symptomatic PAD. Combine that with an Auto-Lookback provision that the RFA strongly encourages, where plans confirm the PA via automated ICD-10 review of health records without even bothering the provider, and what you have is a PA framework that is designed to be permissive, not restrictive. Plans that were hoping to use aggressive utilization management to control GLP-1 volume were being told, explicitly, that they couldn't.

Then you add the formulary uniformity requirements. Every model drug, same tier, same cost share, no disadvantaging any one drug versus another except where FDA labeling forces a variation. This meant a plan couldn't prefer Mounjaro over Zepbound, or Ozempic over Wegovy, or push patients toward orforglipron once it landed. Every model drug had to sit at parity. Combined with the WAC-based gross cost treatment described in Section 2.2.7 (waiving the maximum fair price ceiling for purposes of gross drug cost calculation), the model was asking plans to accept broad, uniform, permissively managed coverage with pricing mechanics they hadn't previously dealt with in Part D.

The pricing itself is where it gets more interesting, because the RFA is coy but not entirely silent. Appendix C lists a $245 net price per month supply, tied at minimum to the Zepbound KwikPen presentations, with TBD markers on several other products and formulations. A $245 net monthly price is, in isolation, a meaningful discount from list and arguably from current plan net cost for commercial lives, but it is not transformative enough to offset the utilization demand shock of opening weight management coverage to a BMI 27 population with qualifying comorbidities. The actuarial question any plan was running was whether the negotiated rebate stack, plus the MDP discount, plus whatever the narrowed risk corridor provided, was enough to cover the volume that would predictably follow from handing out $50 semaglutide and tirzepatide to every qualifying beneficiary. For most plans the answer appears to have been no.

The Bridge as a bridge to what exactly

The Medicare GLP-1 Bridge is a Section 402(a)(1)(A) demonstration, which is a different statutory animal than the Section 1115A authority that powers BALANCE itself. Section 402 is the old Social Security Amendments of 1967 demonstration authority that lets HHS test changes in payment methods that increase efficiency and economy of Medicare services. It's narrower in scope than 1115A but it also requires less procedural overhead, which is why CMS chose it for the Bridge. The Bridge sits outside the Part D benefit and payment flow, which means it doesn't touch plan bids, doesn't touch PDE, doesn't touch the risk corridor, doesn't touch DIR, and doesn't

touch the MDP. It is essentially a direct-to-beneficiary access mechanism, $50 per month for eligible Medicare beneficiaries, running from July 1, 2026 through December 31, 2027 under the extended timeline.

The question is what the Bridge is a bridge to. Originally the answer was clean: the Bridge was a six-month on-ramp to the January 1, 2027 Part D launch of BALANCE. Beneficiaries would get access starting mid-2026 under the demo, then transition into plan-based coverage starting in January. Now that BALANCE is paused, the Bridge runs for a full additional calendar year without any confirmed successor model. The memo's language about additional time and data to inform potential implementation of BALANCE in Part D is carefully noncommittal. Potential means not definite. Data collection means the agency is giving itself optionality.

If you squint at this the right way, the Bridge extension is effectively the policy. The question of whether Medicare covers anti-obesity GLP-1s in 2027 has been answered in the affirmative, just not through Part D plans. It's been answered through a Section 402 demo that sits alongside the program and provides access at a flat $50 to eligible beneficiaries. Plans aren't on the hook. Manufacturers are presumably absorbing the discount delta between demo price and their preferred economics in exchange for volume access to the Medicare population on favorable terms. Beneficiaries get the drugs. Everyone kicks the hard plan-level design problem to CY2028 or later.

There are real limitations to the Bridge approach that will become clearer as the extension plays out. The demo is time-limited, which means manufacturer willingness to keep absorbing the pricing is bounded. The demo doesn't build the formulary infrastructure or the claims data history inside Part D plans that CMS wanted BALANCE to build. The demo also doesn't address the adverse selection question for Part D plans in any durable way, because plans aren't actually carrying the coverage, so the actuarial experience generated is inside the demo rather than inside plan bids. When CY2028 comes around and CMS tries again (assuming it does), plans will still be looking at first-year demand shock without a meaningful internal book of business to price off of. The Bridge buys time but doesn't solve the underlying information asymmetry.

Manufacturer posture, pricing, and the WAC problem

Eli Lilly and Novo Nordisk both agreed to participate in BALANCE, which matters because BALANCE is a voluntary model on both sides. Appendix C lists Zepbound, Mounjaro, orforglipron (pending FDA), Ozempic, Rybelsus, and Wegovy as the covered model drugs, with the net price anchor of $245 per month supply showing up

clearly on Zepbound KwikPen and with TBD markers elsewhere that suggest final negotiations were still in flight as of the March RFA release. Both companies agreed to a "CMS-sponsored model" safe harbor pathway under 42 CFR 1001.952(ii) in lieu of a separate fraud and abuse waiver, and both committed to fund the lifestyle support platform at no cost to beneficiaries or plans.

The interesting mechanical piece is the waiver of Section 1860D-2(d)(1)(D), which says that the negotiated price of a selected drug must be no greater than the maximum fair price plus dispensing fee under the Negotiation Program. Semaglutide, notably, is on the Medicare Drug Price Negotiation Program list, with the negotiated MFP taking effect in 2027. By waiving 1860D-2(d)(1)(D) for purposes of BALANCE, CMS let plans treat gross drug costs based on WAC rather than the MFP ceiling, which materially changes the gross-to-net arithmetic and, importantly, changes how these costs flow through the Part D redesign risk corridor mechanics. This is surgical policymaking. It's also a tacit acknowledgment that layering BALANCE on top of MFP without this waiver would have created weird incentives around reporting and payment reconciliation that no one wanted to deal with.

Now that BALANCE is paused in Part D, both manufacturers are left holding negotiated terms that have no immediate deployment venue in Medicare plan coverage. The Medicaid launches still happen, which gives the agreements partial relevance, and the Bridge presumably operates under some variant of the negotiated pricing architecture, but the big prize (MA-PD and PDP coverage with uniform formulary placement and low cost sharing driving volume) is on ice. For Lilly this is a pause on what would likely have been meaningful Zepbound and Mounjaro Medicare pull-through in 2027. For Novo Nordisk the picture is more complicated because of the MFP overhang on Ozempic and the competitive pressure on Wegovy, but the basic read is the same: negotiated economics that assumed a Part D volume channel now have less of a channel.

Orforglipron is the wild card here. Lilly's oral GLP-1, pending FDA approval, sits in Appendix C with a TBD price. If it gets approved in 2026 or early 2027 (phase 3 data has been strong), the product enters a market where the Part D coverage architecture for anti-obesity indications is still effectively undefined. That is not a disaster for Lilly (the Bridge still provides access, Medicaid still covers, commercial coverage continues to expand), but it meaningfully reduces the near-term Medicare opportunity relative to what a BALANCE launch would have delivered. For a drug whose entire strategic thesis is displacing injectables via oral convenience in a volume-sensitive market, losing the Medicare Part D volume channel for a year is not nothing.

Medicaid keeps rolling and why that matters more than it looks

The memo's confirmation that the Medicaid side of BALANCE proceeds on schedule is probably the most underappreciated piece of this whole story. State Medicaid agencies can apply through July 31, 2026, and can select participation start dates anywhere between May 1, 2026 and January 1, 2027. States that miss the January 1 window are out absent CMS discretion. The BALANCE Model webpage update on April 22 confirmed this timeline.

Here's why this matters more than the Medicare pause suggests. Medicaid populations have historically been excluded from anti-obesity GLP-1 coverage in most states, with a handful of exceptions. The BALANCE negotiated pricing, combined with Section 1115A waivers that let CMS bypass the statutory Medicaid drug rebate mechanics for model purposes, gives state Medicaid agencies an actually affordable way to open up access. States that have been under political pressure to expand obesity coverage but couldn't make the budget math work now have a CMS-negotiated path that materially reduces the per-member cost.

From an investment standpoint, any thesis that was predicated on Medicare Part D as the near-term volume catalyst for Medicare-age GLP-1 demand has to get rewritten, but the Medicaid side is arguably more interesting now because it is proceeding without the 80 percent critical mass problem that killed the Part D launch. States join one at a time, CMS doesn't need a nationwide commitment, and the pricing economics may actually be more favorable for manufacturers on a net basis because Medicaid already has the baseline rebate infrastructure that model pricing gets layered on top of. The watch list here is which states move first. Larger states with recent Medicaid expansion energy and governors who have publicly talked about obesity as a health policy priority are the obvious candidates. The July 31, 2026 application deadline is the near-term action window.

The narrowed risk corridor incentive and why plans shrugged

Section 2.4.1 of the RFA lays out the optional narrowed first risk corridor threshold, which was one of the two major incentives CMS offered plans to offset the participation risk. Under standard Part D risk corridors, the first threshold sits at plus or minus 5 percent of target amount, with plans bearing 100 percent of the variance within that band. BALANCE offered eligible plans the option to narrow that first threshold to plus or minus 2.5 percent, effectively reducing the plan's first-dollar

bid variance exposure by half. The 50/50 and 80/20 corridors beyond the first threshold stayed the same.

The catch, and it's a real catch, is that the narrowed corridor only activates for plans whose model drug utilization rate exceeds one standard deviation above the mean for their plan type among all participating plans that opted in. Appendix D spells out the triggering event methodology: identify participating plans by type (C-SNP, I-SNP, D-SNP, MA-PD, PDP), compute plan-level utilization rates as December enrollees with at least one model drug PDE divided by total December enrollment, calculate mean and standard deviation by plan type, and eligible plans are those with utilization rates more than one standard deviation above the mean.

The problem with this structure is that it only helps plans that get slammed hardest with utilization, and it only helps them on the first 2.5 percent of cost variance. If you're a big PDP with a moderately bad utilization experience (say half a standard deviation above the mean), you get no relief. If you're a plan with truly catastrophic utilization (three standard deviations above the mean), the 2.5 percent relief is a rounding error relative to the actual cost overrun you're carrying. The narrowed corridor is designed to help plans with adverse selection at the margin, not plans that experienced large-scale demand shocks. And adverse selection in anti-obesity coverage for Medicare isn't really a marginal concern. It's potentially structural.

Plans did the math and concluded that the incentive wasn't enough. That's the most honest read of the participation miss. CMS offered a modest risk corridor sweetener, manufacturers offered negotiated pricing, CMS waived a bunch of statutory constraints to make the mechanics work, but at the end of the day the plans that model this stuff for a living looked at the combination and said no thanks. That's a signal worth paying attention to, because it tells you that the next iteration of BALANCE (or whatever its successor is named) needs meaningfully stronger plan-level financial protection, or meaningfully more favorable pricing, or both. The policy didn't fail because the idea was bad. It failed because the economic package didn't clear plan hurdle rates at the level of market coordination required.

340B, the FAD field, and operational plumbing that still has to get built

One of the more technically interesting pieces of the RFA is Section 2.2.8, which addresses the 340B adjustment. Model rebates get adjusted downward by an amount not to exceed 5 percent to account for units purchased by 340B-covered entities at discounted prices. This is CMS acknowledging that manufacturers can't reasonably be expected to pay a full model rebate on top of a 340B discount, which would create

double-discounting in a way that breaks the drug pricing stack. The 5 percent cap is specific, which suggests CMS and manufacturers negotiated hard on exactly how much 340B leakage the model would absorb.

Section 2.2.7 introduces the Facilitated DIR (FAD) field on the PDE, which is a new piece of technical infrastructure specifically for the model. The FAD field calculates manufacturer rebate obligations as the difference between plan-reported ingredient cost on PDE and the sum of GLP-1 Discounted Price plus MDP amount. Manufacturers get invoiced quarterly through the existing Manufacturer Payment Portal but on separate model-specific invoices. Plans don't have to report the FAD rebate on the DIR Report for payment reconciliation, but CMS will incorporate FAD amounts when calculating annual Part D reconciliation.

This is genuinely new plumbing. The PDE format has been modified for this model, the payment portal has been extended, and the reconciliation mechanics have been adjusted. CMS has presumably already built or is building this infrastructure, because the model was supposed to launch January 1, 2027 and you don't spin up PDE field changes in a quarter. That infrastructure now sits idle for a year, which is not the end of the world but it is a real sunk cost that the pause puts in limbo. When CY2028 comes around and CMS tries to relaunch BALANCE (or its successor), the technical pieces should still be largely in place. That's one of the quieter silver linings of the pause: the operational burden of cold-starting the model for a later year is lower than it would have been if CMS hadn't already sunk the infrastructure work.

For RCM vendors, pharmacy IT vendors, and anyone whose roadmap was assuming the FAD field would go live January 1, 2027, the pause is a timing reset but not a cancellation. The spec exists, the field is defined, and once model participation re-emerges, the plumbing is ready. Vendors who were mid-build should not rip out the work. They should probably slow the timeline and reallocate engineering resources to more urgent near-term priorities, but the investment is not stranded.

What the delay means for PBM formulary strategy

PBM formulary teams have had a fun 2026, and the BALANCE pause adds another variable to an already complex model. The CY2027 standalone PDP market was going to be challenging regardless, with the Part D redesign fully absorbed, the $2000 OOP cap creating first-dollar plan exposure, and the MFP round one drugs starting to flow through at negotiated prices. GLP-1 utilization management under BALANCE would have been meaningfully different from commercial book utilization management, with CMS-defined PA criteria that are permissive by historical standards and uniform

formulary placement requirements that prevent the normal lever-pulling on tiering and step therapy.

With BALANCE paused, PBM formulary strategy for CY2027 reverts to status quo ante. Anti-obesity GLP-1 coverage remains not required and largely not provided in Part D. Commercial book coverage continues to expand at the pace employers choose to pay for, which is expanding but not uniformly. Medicaid coverage expands unevenly by state. The Bridge operates outside of plan formulary workflows entirely, so PBMs don't need to integrate Bridge beneficiaries into their formulary tools.

The medium-term question is what PBMs do with their internal planning for a BALANCE 2.0. The big PBMs presumably had workstreams stood up to support plan-sponsor clients in BALANCE participation, including PA adjudication, claims routing, MAC reimbursement logic, and reporting. Those workstreams should probably not be wound down entirely because the model is likely to return in some form. But they need to be rescoped for a delayed timeline, which means headcount decisions, vendor contract decisions, and internal prioritization decisions. PBM CIOs and CFOs are probably having slightly unpleasant conversations this week about capitalized development costs and expected go-live dates.

The more interesting strategic question for PBMs is whether the BALANCE pause creates an opening for private negotiation with manufacturers on GLP-1 pricing for Medicare lives on a plan-by-plan basis. If CMS isn't going to coordinate the market through a multi-lateral demo, nothing in the waiver framework prevents individual plans from negotiating individual rebate deals. Whether manufacturers are willing to engage in that is a different question, and historically they have been reluctant to carve up Medicare channels because of MDP and rebate best-price exposure. But the BALANCE waivers showed that these constraints can be bent for the right demo structure. Sophisticated PBMs are probably already sketching out what a private, bilateral version of BALANCE might look like, using CMS's model architecture as a template.

Read-throughs for investors and operators

For investors in anti-obesity drug manufacturers, the BALANCE pause is a near-term revenue headwind for 2027 Medicare volume, partially offset by continued Bridge access and Medicaid expansion. Consensus models that assumed Part D coverage kicking in January 1, 2027 need to be revised. The magnitude depends on what percentage of 2027 GLP-1 Medicare revenue was going to flow through Part D plan coverage versus the Bridge. Back of the envelope, the Bridge operates at $50

beneficiary cost share with CMS-negotiated manufacturer net pricing, which is broadly similar to the BALANCE net pricing architecture. The real difference is reach and enrollment. Plan-based coverage hits everyone enrolled in a participating plan with minimal friction. Bridge enrollment requires beneficiary action to opt in, which meaningfully reduces effective penetration.

For investors in GLP-1 adjacencies (obesity-related SaaS, lifestyle platforms, nutrition services, remote monitoring, telehealth weight management), the pause is a mixed signal. The lifestyle support platform requirement in BALANCE was going to create real commercial opportunity for platforms that could serve participating plans. That opportunity is delayed. But the underlying demand curve for obesity management services isn't changing. Commercial coverage expansion continues. Medicaid expansion accelerates in certain states. The Bridge creates a large Medicare population with active GLP-1 treatment who need wraparound services. Companies in this space should not over-index on the BALANCE pause as a demand signal. The demand is still there, just flowing through different channels.

For investors in Medicare Advantage plans and PDPs, the BALANCE pause is a near-term positive. Plans avoid the adverse selection risk of opt-in participation, they avoid the bid complexity, and they avoid the PDE format changes that were going to require IT work. CY2027 bids are simpler than they would have been, actuarial assumptions are more stable, and the premium stabilization challenges of the 2027 standalone PDP market don't have a GLP-1 overhang to complicate them. This is a small positive rather than a large one because plans were already going to participate or not on a voluntary basis, but sentiment and bid predictability both improve.

For investors in health IT infrastructure, including pharmacy connectivity, PBM plumbing, revenue cycle management, and drug pricing tools, the pause is a timing reset. The FAD field spec exists, the payment portal extensions exist, the HPMS participation flags exist. None of that infrastructure gets deployed in production for CY2027. Vendors that built toward a January 1 launch are sitting on completed or near-completed features waiting for a use case. The right strategic response is to keep the capability maintained but not to keep burning engineering cycles on enhancement. When BALANCE 2.0 or its successor emerges, the lead time to reactivation should be short.

For entrepreneurs building in the prior authorization, formulary management, and specialty drug access space, the BALANCE pause preserves the status quo for one additional year in Medicare. That is either good or bad depending on the specific product thesis. Products that helped plans manage BALANCE-specific PA under permissive CMS-defined criteria lose a use case. Products that help plans manage

non-BALANCE PA (which is still the reality for anti-obesity GLP-1s in Medicare outside the Bridge) continue to have the same market as before.

What to watch between now and the CY2028 bid cycle

The near-term watch list is specific and manageable. The July 31, 2026 Medicaid application deadline is the first major signal. How many states apply, which states, and what their projected enrollment looks like will tell you how serious the Medicaid leg of BALANCE becomes in the real world. The states that move first will create the early evidence base that CMS uses to argue for a Medicare relaunch.

The Bridge launch on July 1, 2026 is the second major signal. How smoothly the access mechanics work, how quickly beneficiary uptake scales, what the actual utilization curve looks like, and whether there are any surprises in manufacturer supply or pharmacy dispensing all feed into the evidence base for CY2028 planning. Bridge data is exactly what CMS cited in the April 21 memo as the basis for future BALANCE implementation decisions. The quality and completeness of that data matters.

The CY2028 Advance Notice, which typically drops in January or February 2027, is where CMS will signal whether BALANCE is getting a second attempt for plan year 2028 or whether the agency is pivoting to a different architecture entirely. Watch the language around Innovation Center GLP-1 activity in that notice carefully. If BALANCE language reappears in substantive form, the relaunch is on track. If the language softens to generic references about continuing to evaluate demonstration options, the relaunch is at risk of being delayed further or restructured.

The FY2027 appropriations fight will matter too. CMMI's authority is broad under Section 1115A, but political dynamics around obesity coverage, drug pricing, and Medicare spending have intensified, and Congress has shown increasing willingness to weigh in on specific Innovation Center models through hearings, letters, and appropriations rider discussions. Any signal from key committees about BALANCE specifically or about anti-obesity Medicare coverage generally is worth tracking.

Finally, the CY2028 Part D bid cycle itself, which runs through spring and summer 2027, is where the next practical test of plan willingness comes. If CMS tries to relaunch BALANCE for CY2028 with meaningfully stronger incentives (narrower risk corridors across the board, not just for top-quartile utilizers, or higher guaranteed rebate stacks, or some form of direct federal backstop), plans will look at the package and decide again. The 80 percent threshold stays a hard constraint unless CMS lowers

it, which would be a significant policy concession. Watch whether the threshold gets revisited. That single number tells you most of what you need to know about whether the next attempt is designed to succeed or designed to be a second attempt with similar odds of missing.

The BALANCE pause is not the end of Medicare coverage for anti-obesity GLP-1s. It is an iteration in a multi-year design process that is clearly still underway. The Bridge extension buys the agency time, preserves beneficiary access, keeps the manufacturer negotiations warm, and generates the utilization data that will inform the next attempt. Plans got a reprieve on a risk they didn't want to take this year. Manufacturers lost a Medicare Part D volume channel but kept the Medicaid leg and the Bridge. Investors in the space had to revise their 2027 Medicare volume assumptions but the underlying demand thesis is unchanged. The interesting question, for everyone who has been tracking this closely, is whether the next attempt looks fundamentally like BALANCE with better incentives, or whether CMS comes back in a year with a meaningfully different architecture. That answer starts coming into focus in the CY2028 Advance Notice, and it fully reveals itself in the CY2028 bid results next summer.

← Previous

Discussion about this post

Comments Restacks

Write a comment...