Turquoise

PUBLIC COMMENT

CMS-1850-P
RIN 0938-AV83

Medicare Program: Hospital Outpatient Prospective Payment and Ambulatory Surgical Center Payment Systems; and Quality Reporting Programs; Including the Hospital Outpatient Quality Reporting Program and Ambulatory Surgical Center Quality Program; Request for Information on Strengthening the Standardization and Comparability of Hospital Price Transparency (HPT) Data; Prior Authorization; Accrediting Organization (AO) Deeming for Emergency Medical Treatment and Labor Act (EMTALA); and Notices of Closure of Teaching Hospitals and Opportunities To Apply for Available Slots

Comments prepared by Turquoise Health

August 2026

Turquoise

Introduction

Turquoise Health (Turquoise) thanks the Centers for Medicare & Medicaid Services (CMS) for the opportunity to publicly comment on hospital price transparency machine-readable files (MRFs), shoppable services standard charge information, and the 340B program, all published as part of the CY 2027 OPPS/ASC Proposed Rule.

We focus our comments on areas and examples of what we have seen work, what we have observed hospitals and payers struggle to report consistently, and where we believe targeted structural changes would meaningfully improve the usability of this data for patients, researchers, employers, and innovators alike.

Summary

Regarding MRFs and reporting shoppable services, the Turquoise team draws on our experience aggregating and parsing MRFs at scale to recommend that

1. CMS move toward a structured, object-oriented schema for contract terms such as outlier, stop-loss, and carve-out provisions, and
2. Consumer-facing price disclosures are built around bundled, comprehensible service packages rather than raw codes.

Turquoise fully supports CMS' plans to include stop-loss/outlier rates in MRFs. Stop-loss represents a significant portion of the cost of healthcare, and without access to the negotiated rates, the industry is unable to get a full picture of cost estimates.

Regarding the proposed 340B payment adjustment, we use our own market data and a focused high-cost oncology drug example to show that the proposed rate change as written may affect large systems and safety-net hospitals differently. We flag two interactions the rule does not currently address:

1. the transition treatment for drugs exiting pass-through status, and
2. the unresolved interplay with the Inflation Reduction Act's Maximum Fair Price program.

Across all topics - MRFs, shoppable services, and 340B payments - our recommendations are generally focused on giving existing logic clearer structure rather than asking the industry or CMS to invent and publish something entirely new.

Turquoise

Turquoise Responses to the Proposed Rule RFI

Machine-Readable Files

1. Increasing Transparency of Outlier Provisions and Additional Contract Terms

What information is needed for users of the MRF to fully understand contract terms related to outlier payments, stop-loss, rate-tiering, carve-outs, and other adjustments?

An MRF user typically needs two things:

1. the actual parameters behind each contract term, which vary meaningfully, and
2. an unambiguous, machine-readable link between those parameters and the specific rates they modify.

That second category is where the CSV format falls short. CSV is row-based, meaning every row stands alone, with no reliable way to reference a shared object defined elsewhere in the file. In practice, hospitals with terms that span thousands of code and payer combinations respond in one of three ways:

1. duplicating the full contract language on every affected row, which causes significant file bloat
2. collapsing the term into a vague flag that strips out the logic that actually determines payment
3. omitting it entirely.

To most effectively aggregate and parse reported contract terms and adjustments, **we recommend CMS move toward a JSON-based, object-oriented schema.** JSON allows a hospital to define a contract-term object once, assign it an identifier, and reference that identifier across every rate to which it does or does not apply without duplicating or losing logic. JSON also supports term-type-specific parameters, since the reporting elements relevant to a stop-loss provision (a threshold and a conversion rate) are not the same as those relevant to a lesser-of or not-to-exceed provision.

We recognize that JSON is less spreadsheet-friendly than CSV, but multi-gigabyte CSVs with unstructured free text scattered throughout are not meaningfully more usable today. A computer can more reliably parse a JSON, and reliably parsing MRFs at scale is foundational to the utility of the data at large.

How are hospitals currently reporting contract terms for outlier payments, stop-loss, rate-tiering, carve-outs, or other adjustments? Are current formats sufficient?

Current reporting is inconsistent and, in our experience monitoring files across the industry, largely ineffective. We have seen stop-loss and outlier terms posted in additional-notes fields, algorithm notes, generic notes, and general contract provisions objects, often for the same term type, across different hospitals or even across different files from the same hospital. This approach makes files cumbersome, given the challenging nature of reporting a block of text within a file. The examples below typify various lengthy stop-loss definitions that may be seen in files:

> Inpatient services over $1,000,000.00 have a tiered reimbursement for all Insurance Plans. When total claim charges exceed $1,000,000.00 to $1,500,000.00, charges

exceeding $1,000,000.00 will be reimbursed at 90% of charges. This amount will supersede any other contracted rate that would otherwise apply. When total claim charges exceed $1,500,000.00, charges exceeding $1,500,000.00 will be reimbursed at 80% of charges. This amount will supersede any other contracted rate that would otherwise apply. Clean Claims paid within fifteen (15) calendar days of receipt by Insurance shall be eligible for an additional 5% discount - 73% billed charges. This prompt pay discount is only applicable to hospital services. Inpatient services over $1,000,000.00 have a tiered reimbursement for Insurance plan: 95% up to the first $1,000,000.00 of covered charges; 90% of the next covered charges from $1,000,000.00 to $1,499,999.99 and 80% of any charges over $1,500,000.00. These amounts will supersede any other contracted rate that would otherwise apply.

—

Stop-loss formulas may be assigned at the category or global levels. Should a contract require stop-losses to be assigned at both the category and global levels, the category level stop-loss will take precedence for those accounts matching both the category and global stop-loss criteria. Just as certain parameters can be set up to be excluded from the stop-loss calculation. If the basis is set up as a Standard Reimbursement threshold, the system will calculate the expected standard reimbursement amount and compare it to the threshold amount in the stop-loss formula. If the Standard Reimbursement amount exceeds the stop-loss threshold, the payment is based on the % or flat amount specified in the stop-loss formula instead of the Standard Reimbursement. If the basis is set up as a Charges threshold, the system will determine if actual charges exceed the stop-loss threshold. If so, the payment is based on the % or flat amount specified in the stop-loss formula instead of the Standard Reimbursement. The system will calculate the reimbursement amount indicated in the contract and then will compare the resulting figure to the percent of total charges indicated in the stop-loss. The greater of the two amounts will then be compared to the minimum and maximum per case rates to indicate the expected reimbursement. (The minimum and maximum fields are optional based on the contract language.) The system will calculate the reimbursement amount indicated in the contract and then will compare the resulting figure to the percent of total charges indicated in the stop-loss. The lesser of the two figures will then be compared to the minimum and maximum per case rates to determine the expected reimbursement. (Minimum and Maximum case rates are optional.) The system will calculate the reimbursement indicated in the contract and then will compare the result to the minimum and maximum percent results and/or amounts indicated in the stop-loss. If the calculated amount falls between the minimum and maximum amounts, that figure will be the account's expected reimbursement. If the calculated amount falls below the minimum, the minimum amount will be considered the expected reimbursement. If the calculated amount is greater than the maximum amount, the maximum will be the expected reimbursement. This stop-loss formula provides the ability to stop the reimbursement for an outpatient account once the amount reaches the reimbursement amount for the DRG if the account were an IP account. This is applied for OP formulas at the category level.

Recent CMS guidance narrowed the ambiguity somewhat, but we continue to see hospitals default to treating every inpatient rate as an algorithm simply because a stop-loss clause exists in the contract. That approach may be technically defensible, but it is analytically counterproductive because it obscures the actual applicable rate rather than revealing it. **Simply put, treating every inpatient rate as an algorithm to account for stop-loss is not a useful model. CMS should not implicitly endorse it by declining to standardize further.**

Do outlier and carve-out contract terms typically apply to individual services, a category of services, or across the entire contract? Should rate-tiering be reflected in the payer-specific negotiated charge?

In our experience, it depends on the terms. Any standard needs to accommodate both ends of the services spectrum rather than force a single model. Certain provisions, such as stop-loss and outlier language specific to transplants, CAR-T, and other high-cost service lines, are highly specific and effectively service-level. Others, such as the most common “first-dollar” inpatient stop-loss, apply broadly across an entire contract or claim type (frequently DRG-based), sometimes with categorical exclusions for separately paid drugs and implants, or for service lines like NICU, maternity, or rehabilitation that are modeled independently. **We would characterize a stop-loss provision that carries such exclusions as a general contract provision noting an exception, not a service-specific term.** CMS's standard should support both a general_contract_provisions-style object for broadly applicable terms and a direct code-level linkage for the narrower ones. The service-specific model will apply less often than the broadly applicable one, but both need a defined home within MRFs.

We note that a JSON schema can more naturally support this type relationship where 1 to n number of provisions can apply to a single rate.

What additional contracting or payment adjustments affect payer-specific negotiated charges, and should that information be encoded in current data elements?

As mentioned above, we reiterate our stance that any contract language capable of altering the negotiated rate should have a defined home in the schema. That should not be limited to not only stop-loss and outlier provisions, but lesser-of and not-to-exceed clauses, bundling and multi-procedure reductions, and transfer-related adjustments. These are the same terms used in claims pricing today that have no consistent home in the MRF. Thus, hospitals either invent their own free-text convention for each one or omit the logic altogether. The lack of definition within the schema limits the utility of each of these contract elements, and all are crucial to understanding payer-specific negotiated charges in the context of estimate tools and distilling down how much an item or service will ultimately cost a patient.

A JSON schema with term-type-specific object definitions, as recommended above, is the most direct way to give each of these adjustments a governed structure rather than another ungoverned free-text field. We note that this same design has direct precedent in the Transparency in Coverage (TiC) files, where CMS's CSTM-ALL billing code type already establishes a mechanism for record-level disclosure of terms, including outlier-type provisions, that do not need to be tied to a specific billing code. Extending an analogous approach to the hospital MRF, rather than inventing a wholly separate model, would minimize the burden on parties that report both files.

2. Standardization to Enhance Utility of the MRF

Are there challenges in parsing and categorizing free text fields? What would facilitate parsing?

Turquoise observes a number of challenges in parsing and categorizing free text fields, some of which we’ve touched on in previous answers.

At minimum, free-text notes should carry a required tag or category identifying what the note refers to so that a stop-loss string is unambiguously labeled as a stop-loss note rather than left to inference. More fundamentally, the parsing burden is a symptom of the format problem

described above: a schema with term-type-specific objects and parameters would eliminate most of the need to parse free text at all, because the underlying information would already be structured rather than narrated. We observe other industry organizations and groups have created detailed, prescriptive schema-based solutions to a more organized approach to the information currently located in a free text field that CMS may benefit from reviewing and considering.

Are there categories of information consistently included in free text fields that could be standardized through new data elements?

Yes. In addition to the contract-term categories described above we also ask CMS to consider rate-derivation logic. We frequently see free-text notes used to explain that a rate is derived from a fee schedule or a base rate and weight. Common examples include reported rates such as “150% of Medicare OPPS” or “Base Rate of $10,000.” This is exactly the kind of information that should be captured as a structured data element rather than left as a string a downstream user has to interpret correctly.

For additional non-code based terms that may impact reimbursement such as age, length of stay, discharge status, bill type, and others, CMS should consider providing guidance regarding if the use of CSTM-ALL is correct for reporting these terms or if CMS intended hospitals to employ a different approach to capture and report them.

Are there additional contract methodologies that should be reflected as valid values in the MRF, beyond fee schedule, capitation, per diem, case rate, and other?

Yes. In a similar way that TiC updates are geared towards a seamless ability to link and compare hospital and payer MRFs, the hospital final rules should aim to resolve hospital file inconsistencies relative to payer TiC files.

For example, percent of charge is a common commercial contracting methodology that is not currently a valid value for hospital MRFs. CMS should also support combination or contingency methodology types, such as a case rate that converts to a per diem after a defined threshold, because these methodology types are common in the market but not representable under a single-methodology model today. This overlaps meaningfully with the algorithm concept CMS has previously solicited feedback on. We would encourage CMS to expand valid rate and methodology types as part of that same effort rather than as a separate exercise.

What additional information should CMS consider to enhance the utility and comparability of MRF data?

Every rate should carry an effective date and expiration date tied to the specific contract term that produced it. Assuming hospitals remain on an annual MRF update cadence, a hospital could in theory publish its file days before a broad rate refresh takes effect. Without effective-date metadata at the rate level, a user of the file has no way to know whether the rates in front of them are current or about to lapse. This predicament undermines the comparability goals CMS is trying to accomplish.

Shoppable Services

As a primer to the following comments, Turquoise has built its Standard Service Package (SSP) framework specifically to address the gap CMS identifies between raw MRF data and a price a patient can actually use. SSPs bundle the codes, ancillary services, and fees associated with a procedure into a single consumer-facing package, and assign each one a Shoppability Index to

identify (A) how easily a patient can shop for the service and (B) the degree to which price is a key lever in the decision. We share that framework and the operational experience behind it in the responses below.

Should CMS revisit the number and types of CMS-specified shoppable services? What are the advantages and disadvantages of increasing or decreasing the list?

At Turquoise, we look at shoppability as a function of a patient's willingness and ability to browse providers, find a price, and act on it. The most shoppable services tend to be non-emergent and low in specialization, and we assess them across several factors:

- Plannability: can the service be scheduled in advance?
- Optionality: are comparable providers available?
- Price variability: is there real potential for savings?
- Outcome predictability: is the service defined enough that its scope and resources are known before it begins?

We believe CMS would benefit from using a similar assessment to revisit the types of shoppable services from the original final rule. Rather than simply expanding a curated list of 70 items, we would recommend **CMS base shoppable services on a consistent, quantitative measure of shoppability that can be reassessed periodically, rather than an administratively fixed set.** The advantages of moving away from a fixed list of shoppable services are that CMS, innovators, and patients are more frequently given access to understand the cost of the care they'd be most likely to review on the factors listed above. The disadvantage is the challenge and responsibility to maintain an updated list of services.

We would also encourage CMS to classify each required service into one of three buckets:

- a single item or service
- an encounter
- an episode

These buckets allow consumers to understand whether a given estimate reflects one line item or a broader bundle of care. Where a code represents an encounter or episode, the associated ancillary components should either be required in the estimate or explicitly disclosed to the patient as excluded, so that the scope of the stated price is clear and consumers have a standardized basis for comparison across providers.

What are the advantages and disadvantages of requiring a shoppable services file, versus removing deemed compliance for price estimator tools? Does the current split make cross-hospital comparison harder?

We clarify if the advantages or disadvantages stated in this question are for **CMS and innovators downloading files en masse** or **patients seeking to understand the cost of a specific treatment.** In our experience, **patients are not visiting provider websites to download MRF or shoppable service files.** They are either accessing a patient estimate tool or working with their insurance provider to get an understanding of how much a service will cost based on their plan details. Before patients even enter any sort of cross-hospital comparison exercise, they first need to meaningfully engage with a tool of some sort that lays out costs in plain language. In that scenario, it's our opinion that a shoppable services file is an insufficient solution for patients to rely on.

However, data available in a common format or schema is always beneficial to any entity looking to surface a subset of rates in some sort of estimate tool or conversational AI interface. Assuming all shoppable services are accurately reported in an MRF, which they should be, in accordance with the previous hospital final rules, Turquoise's bigger area of focus is ensuring the right codes and rates get bundled and shared with patients in a way that patients do not engage with MRFs or files at all.

We observe that the current split does make cross-hospital comparison harder. The core problem with deemed compliance today is that it allows a hospital to satisfy the consumer-friendly display requirement without a clear data trail. An MRF can be checked against the CMS validator tool line by line, but a web-based estimator's output cannot be compared the same way, because there is no guarantee an MRF and a PET compute or disclose their figures using the exact same calculations or reporting structure. Both the MRF and the PET may be fully compliant; however, as mentioned in our earlier paragraph, the final output may differ because the solutions are built for different audiences.

For example, some hospitals attempt to bundle all codes into an episode of care within an estimator tool, and those bundled items and services are necessarily included as individual line items with negotiated rates in an MRF. The result is a two-track system: hospitals that publish files are directly comparable to one another, and hospitals that rely on an estimator tool are not comparable in any systematic way.

Therefore, if CMS retains a deemed compliance path for estimator tools, we believe it should be conditioned on the tool exposing its underlying line-item data in that same structured format, preserving the convenience of the consumer tool while restoring the comparability that the file format provides. We provide examples of how to execute this type of consumer tool further down in our comment.

For hospitals that satisfy consumer-friendly display through a price estimator tool, what mechanisms could make the underlying data available in a separate file?

We recommend requiring an accessible API or bulk export from the estimator tool, structured to the same schema as the shoppable services file requirement. This would allow the same downstream tools and consumer applications, including conversational and AI-based patient tools that continue to emerge, to consume estimator tool data on equal footing with data derived from a full MRF. That data includes shoppable services rates, regardless of which shoppable services compliance path a given hospital chooses.

Are there circumstances in which the standard charges in a shoppable services file would not match the MRF? Why might a hospital disclose a discounted cash price in an estimator tool but not in its MRF?

In our experience, this divergence typically arises when a hospital's finalized discounted cash price is determined internally after care occurs as a discretionary, case-by-case financial assistance or self-pay accommodation rather than a formally adopted pricing policy. That does not necessarily mean the price quoted to a given patient through the estimator tool is illegitimate, though, because patients may be accessing estimator tools before they have provided any necessary paperwork to know which level of financial assistance they qualify for.

Because the discounted cash price data element in the MRF is reporting a formal, executive-approved policy, hospitals with estimator tools built on separate self-pay or financial-assistance logic may not necessarily route back into an MRF.

Which data elements are important for consumers to compare hospitals? Would a standard shoppable services file template help? How should the file indicate ancillary services included or excluded?

As mentioned in our previous answer, we have not observed patients engaging with shoppable services files in a meaningful way. However, a standard template built around bundled service packages, rather than singular codes, would materially improve comparability because it removes the burden on consumers to know which ancillary codes to add to a base procedure code to understand their real cost. In our own SSP framework, each package includes the following:

- consumer-facing name and description
- service and disease area categorization
- shoppability classification
- aggregate view of the ancillary services, codes, and fees typically billed alongside the primary procedure

A standard template with equivalent fields showing a single bundled price, a defined and consistently labeled set of inclusions, and explicit callouts for common exclusions would let consumers compare a hospital's price for “a knee replacement” rather than needing to know “Major hip and knee joint replacement without Major Complications or Comorbidities” or having to try to reconstruct that same price from a dozen separate line items across two different files that require the patients to have clinical coding expertise.

How should hospitals present ancillary items, implants, and bundled services so consumers understand total expected costs? What approaches best distinguish included versus excluded services?

Building on our previous answer, we recommend an open license all-in package price as the primary figure, with an itemized breakdown available beneath it showing what is included in that price, plus explicit, consistently formatted callouts for the categories most often excluded in practice, such as implants, drugs, anesthesia, and physician professional fees. Consumers are best served by a single number they can trust as close to complete, paired with a clear, short list of what is not yet in that number, rather than a complete itemization they are expected to interpret themselves.

What additional information should CMS consider to enhance the comparability of consumer-friendly display data between hospitals?

Two pieces of information would significantly move the needle:

1. Open license standardized service groupings to ensure a displayed service is actually the same clinical service across hospitals, rather than each hospital defining its own bundle informally.
2. Format parity with the MRF itself, so that the tools built to process full MRFs can also process shoppable services data without a separate parser.

In our experience, comparability is ultimately a function of consistent definitions and consistent formats. The same principles that should guide CMS's MRF standardization work apply to shoppable service files.

Turquoise

Turquoise Public Comments regarding 340B

Survey Methodology (Section V.B.8)

1. Trimming of anomalous acquisition-cost records

We seek comment on our proposal to trim records with per-unit acquisition costs more than three standard deviations from either the 340B or non-340B geometric mean per-unit acquisition cost for that NDC.

We have no basis to challenge the specific three standard deviation cutoff, but we note that the resulting acquisition-cost margin of ASP - 33.4% is the single number the entire payment policy rests on. Thus, the sensitivity of that margin to the trimming threshold has significant downstream impact.

We ask that prior to finalizing this proposed rule, CMS consider publishing the acquisition cost distribution before and after trimming, by hospital type, so commenters can evaluate whether the trimmed margin is representative of the DSH hospitals most affected by the policy. Visibility into that distribution would allow the industry to determine what impact, if any, large volume purchasers are having on hospitals treating underserved populations.

2. Approach to identifying and excluding anomalous data

We solicit comment on our approach to identifying and excluding anomalous data, as well as on the overall methodology used to evaluate the survey data for purposes of this proposal.

We observe that market data corroborates CMS's general finding that a substantial gap exists between ASP + 6% payment and 340B acquisition cost. However, an Illinois-specific analysis found that 340B profit and reliance are unevenly distributed across hospital size and type, with large academic medical centers and regional systems capturing a disproportionate share of 340B margin relative to small and rural DSH hospitals.[1]

We recommend that CMS evaluate their anomalous data methodology and report results separately for large multi-hospital systems versus independent and small DSH hospitals. As noted in our previous response, the risk of using a single national margin to make policy decisions is that the value obscures how differently, if at all, this policy will impact hospitals with different profiles.

Proposed 340B Payment Methodology (Section V.C)

Illustrative Example: Per-Dose Margin Impact on a High-Volume Oncology Drug (Pembrolizumab / Keytruda, HCPCS J9271)

Before responding to the specific requests for comment in this section, we want to illustrate the payment methodology's per-drug, per-payer-mix effect using publicly available pricing inputs for pembrolizumab (Keytruda). Keytruda is a high-volume, high-cost oncology drug and one of the largest categories of 340B drug spending nationally. The figures below are Turquoise's own illustrative modeling, using the Q3 2026 ASP for J9271 and representative acquisition cost and

[1]Dan Snow, “The 340B Program Has Gone Off the Rails: An Illinois Case Study of One Federal Program’s Unchecked Growth and Unintended Consequences,” Price Points (Turquoise Health), July 17, 2025. https://www.pricepoints.health/p/il-340b

commercial rate assumptions. We believe the figures usefully visualize the aggregate percentages in the proposed rule into a per-dose, per-hospital-type comparison that the rule's current impact analysis does not provide.

Scenario*	340B Medicare margin	340B commercial margin	Non-340B Medicare margin	Non-340B commercial margin	340B margin break-even Medicare mix
Before proposed rule (both at ASP+6%)	+$2,500	+$11,950	-250	+$9,200	340B CE: Always positive Non-340B CE: 97% Medicare
After ruling, commercial unchanged	–$2,054	+$11,950	-250	+$9,200	340B CE: 85.3% Medicare Non-340B CE: 97% Medicare
After ruling, commercial also cut 30% (irrespective of 340B status)	–$2,054	+$5,440	-250	+$2,690	340B CE: 72.6% Medicare Non-340B CE: 91% Medicare

***Assumptions:** ASP (as of April 2026) $57.78/mg: $11,556/dose at 200 mg | Medicare non-340B (ASP+6%) = $12,250 | Medicare 340B proposed (ASP−33.4%) = $7,696 | Commercial base rate = $21,700/dose · 340B acquisition cost = $9,750/dose | Non-340B acquisition cost = $12,500/dose

Two things stand out.

1. The proposed rule flips the 340B covered entity's Medicare margin on this drug from positive ($2,500/dose) to negative (-$2,054/dose). That swing is exactly equal to the $4,554 gap between ASP + 6% and ASP - 33.4% on this HCPCS code. The non-340B hospital's Medicare margin (-$250/dose) is unaffected and remains negative throughout. **Thus, for this drug, the Medicare-side economics of 340B participation reverses relative to non-participation, for the second time since the program's creation.**

2. Once the Medicare margin is negative, a 340B hospital can only stay whole on this drug by leaning on commercial volume. The break-even Medicare mix column demonstrates a path to staying whole, but it's built on an assumption that commercial payers are reimbursing at ~200% of ASP. With commercial margin unchanged, the hospital can tolerate Medicare making up no more than 85.3% of its volume for this drug before it becomes a net loss overall. If commercial payers respond to the Medicare cut by compressing their own rates by even 30%, which is a foreseeable outcome given how frequently commercial contracts reference Medicare or ASP-based benchmarks, the tolerance falls to 72.6%.

We include this detailed example and data analysis in our comment because it shows that the changes in the proposed rule currently may be insufficient to protect the hospitals likely to be hit the hardest. A hospital's Medicare share of volume is structurally higher at safety net and rural DSH facilities, which by definition serve a disproportionate share of low income and elderly patients relative to commercially-insured patients.

Similar to our comments above, we recommend CMS avoids relying solely on aggregate national dollar impact as is currently presented in the final rule. Instead, CMS should model a per-drug, per-payer-mix effect explicitly, for a sample of high-cost, high-Medicare-utilization

drugs. This approach will allow CMS to evaluate the risk that commercial payers compress 340B-related commercial rates in response to this proposal, because if they choose to do so, hospitals that rely on 340B would continue to see their subsidies erode and may discontinue to provide treatments to patients with high unmet needs.

1. Grouping hospitals by 340B covered entity status (V.C.2)

We seek comment on these proposals. Specifically, we seek comment on our proposal to use 340B covered entity status as a relevant characteristic to group hospitals for purposes of payment based on average acquisition cost under section 1833(t)(14)(A)(iii)(I) of the Act and our proposal to vary the amount of payment for 340B-acquired drugs by the group of hospitals that are enrolled in the 340B Program to more appropriately align payment with the average acquisition cost for the drug.

We agree that 340B covered entity status is a reasonable grouping characteristic under section 1833(t)(14)(A)(iii)(I). Our concern is not with the grouping variable itself but with treating all 340B-enrolled hospitals as financially homogeneous once grouped. The Illinois analysis referenced in a previous response found that outpatient drug revenue makes up an increasing and outsized share of net patient revenue at large DSH hospitals (nearly 20%) while small and Critical Access Hospitals collectively generate far less 340B-derived profit than a single large academic medical center.

A single uniform rate applied to this entire group will reduce more absolute dollars at large systems but will remove a proportionally larger share of available margin at smaller, less commercially competitive hospitals.[2]

Separately, analysis of over 93 million negotiated commercial rates found that for rural hospitals, 340B margin is often one of few available levers to offset already-low commercial reimbursement. As the Keytruda example above illustrates, there are material financial impacts for even just one high-volume oncology drug if the proposed rate turns a positive Medicare margin negative and requires the hospital to monitor Medicare volume just to break even. Hospitals with a structurally higher Medicare share of volume, which is a defining characteristic of the safety net hospitals the 340B and DSH programs are meant to serve, are the ones most exposed to this dynamic.

We recommend CMS supplement the proposed grouping with an analysis of impact by hospital size and rural/urban status, not solely by 340B enrollment status.[3]

2. Periodic revalidation of the survey (V.C.6)

We seek comment on this proposal.

Turquoise supports periodic revalidation and encourages CMS to make the survey responses publicly available. We would also recommend CMS commit to a specific interval rather than leaving the cadence open-ended, given that 340B enrollment and drug mix both change quickly. The rule currently suggests a loose “perhaps as often as every 4 years” revalidation window.

---

[2]Dan Snow, “The 340B Program Has Gone Off the Rails: An Illinois Case Study of One Federal Program’s Unchecked Growth and Unintended Consequences,” Price Points (Turquoise Health), July 17, 2025. https://www.pricepoints.health/p/il-340b

[3]Dan Snow, “Rural Hospitals Are Paid Less Than Urban Hospitals: An Unprecedented Look at the Urban-Rural Reimbursement Gap Using 93M+ Negotiated Commercial Rates,” Price Points (Turquoise Health), Nov. 6, 2025. https://www.pricepoints.health/p/rural-v-urban

Additional annual or 2 year revalidation may be helpful to assess 340B discounts. These discounts are by default 23.1% but they can increase for various factors such as price increases above CPI-U or breaking "best price" by the manufacturer. A consistent and reliable revalidation approach helps prevent any pullback from actions that may cause 340B prices to deteriorate faster to keep reimbursement afloat.

3. Exemption of children's hospitals, PPS-exempt cancer hospitals, and rural SCHs (V.C.8)

We seek comment on our proposal to exempt children's hospitals, PPS-exempt cancer hospitals and rural SCHs from the 340B drug payment adjustment.

We support these exemptions. In addition, we note that a separate analysis of the One Big Beautiful Bill Act's effect on Medicaid spending found that over 300 hospitals, largely disproportionately rural and mid-sized urban DSH hospitals and not necessarily rural SCHs specifically, are at risk of falling below the DSH percentage threshold required for 340B eligibility altogether as Medicaid patient days decline.[4] Therefore, we recommend CMS coordinate the final exemption list with the DSH-threshold eligibility and consider extending exemption or some sort of relief to DSH hospitals near the eligibility threshold.

340B Considerations Not Currently Addressed in the Proposed Rule

The two issues below are not the subject of an existing request for comment in the proposed rule. We raise them because we believe the final rule should address them explicitly, given their tie in to the topics in the proposed rule.

Transition treatment for drugs exiting pass-through status

The proposed rule exempts pass-through status drugs (status indicator G) from the 33.4% reduction, reporting them with modifier TB rather than JG. Pass-through status is time-limited and typically runs two to three years post-launch. The proposed rule does not address what happens at the moment pass-through status expires. As currently drafted, we read the proposal to mean a drug moves overnight from full exclusion to the full 33.4% reduction with no transition.

Given that HRSA's own aggregate purchase data shows the largest 340B dollar volume concentrated in high-cost specialty drugs for conditions like cancer and HIV, the same drug classes most likely to launch with pass-through status, we foresee the potential for a reimbursement cliff at precisely the point when prescribing volume for a new therapy is typically increasing post-launch. We recommend CMS propose a phased step-down (for example, over two to four quarters) for drugs exiting pass-through status at 340B covered entities, and solicit public comment on that transition design in the final rule.

The pembrolizumab example in Section II above illustrates the magnitude of what such a cliff could look like: a $4,554-per-dose swing in Medicare margin for a single high-cost drug. We would expect a comparable, if not larger, swing for a newly launched specialty drug transitioning out of pass-through status, since such drugs typically carry even higher list prices in their first years on the market.

Interaction with the Inflation Reduction Act Medicare Drug Price Negotiation Program (Maximum Fair Price)

[4]Dan Snow, "The OBBB May Disqualify Hundreds of Hospitals from the 340B Program," Price Points (Turquoise Health), July 24, 2025. https://www.pricepoints.health/p/obbb-340b

We reviewed the proposed rule in full and found no reference to Maximum Fair Price, MFP-selected drugs, or the Medicare Drug Price Negotiation Program anywhere in the discussion of the proposed 340B payment policy. As additional Part B-relevant drugs enter the negotiation program in future selection cycles, we believe this interaction needs to be addressed in the final rule rather than left to sub-regulatory guidance. Specifically, we ask CMS to clarify and solicit comment on:

- For a drug that is both 340B-acquired and MFP-selected, whether the proposed 33.4% reduction would be applied against ASP, against MFP, or against the lesser of the two.
- How this proposed 340B adjustment would interact with the payment methodology and add-on already established for MFP drugs under Part B, and whether the two policies could stack in a way that pushes effective reimbursement below a hospital's actual acquisition cost, rather than simply aligning payment with acquisition cost as intended.
- Whether CMS has modeled the combined effect of MFP and the proposed 340B adjustment on hospitals that rely on high-cost, MFP-eligible specialty and oncology drugs to cross-subsidize other service lines.

Related to both issues above, we observe that the combined uncertainty of an eventual pass-through cliff and an unresolved MFP interaction could slow adoption of newly approved or newly launched drugs at 340B covered entities with a heavy Medicare population at precisely the point in a drug's lifecycle when broad access matters most for patients. We recommend the final rule's regulatory impact analysis explicitly discuss this potential effect on adoption.

Conclusion

CMS's price transparency framework has already produced real, measurable data at scale. The gap that remains is largely structural: a file format that cannot cleanly represent the contract logic that actually determines what gets paid, and a consumer-facing display requirement that permits two incompatible paths to compliance. We believe the recommendations above — a move toward a structured, object-oriented MRF schema for contract terms, and a shoppable services standard built around bundled, consumer-comprehensible packages rather than raw codes — would close that gap without materially increasing the reporting burden on hospitals, because in most cases they formalize logic hospitals already apply internally rather than asking for new logic to be invented. We appreciate the opportunity to comment and welcome the opportunity to discuss any of the above in further detail.

Chris Severn - CEO and Cofounder

Carol Skenes - Chief of Staff

Ryan Kost - Senior Director, Data Product

Ariela Simerman - VP of Product

Mischka Moechtar - Senior Manager, Life Sciences

turquoise.health