

# GUARD and GLOBE Walk Into a Spreadsheet: How CMS Accidentally Created the Next Health Tech Infrastructure Cycle

JAN 05, 2026



Share

---

## Abstract

GUARD and GLOBE represent CMS dragging pharmaceutical pricing into the cage through brute force regulation. GUARD tackles Part D drugs, GLOBE targets Part B drugs, but both share DNA: mandatory participation, international reference pricing benchmarks, quarterly reconciliation hell, and a ten-day suggestion of error window that makes tax season look relaxed. The proposals create new infrastructure demands around benchmark replication, rebate forecasting, invoice validation, coinsurance adjustment mechanics, and audit defense. Winners won't be dashboard. Winners will be boring systems that become source of truth for model eligibility determination, unit normalization across countries, PPP adjustments, accrual-governance forecasting, and reconciliation workflows that don't melt under Federal Register scrutiny.

## Table of Contents

What CMS Actually Did and Why It Matters

The Mechanical Reality of Both Models

Why These Models Create Actual Markets Instead of Consulting Theater

Business Model One: Benchmark Replication as Mandatory Infrastructure

Business Model Two: Rebate Accrual Systems That Survive Quarter Close

Business Model Three: Provider-Side Coinsurance Correctness Tools

Business Model Four: Unit Normalization as Boring Infrastructure

Business Model Five: Shortage Intelligence and Supply Chain Disruption Tracking

Business Model Six: International Data Submission Management

Technical Architecture That Passes Federal Audit

Where This Breaks in Practice and How to Build Around It

Why Boring Infrastructure Wins When Policy Gets Specific

## **What CMS Actually Did and Why It Matters**

On December 23, 2025, CMS dropped two proposed rules that most health tech misread as “international drug pricing policy” and move on. That would be wrong. GUARD (Guarding U.S. Medicare Against Rising Drug Costs) and GLOBE (Global Benchmark for Efficient Drug Pricing) are infrastructure specs disguised as payment policy. They mandate specific computational workflows, define data exchange formats, establish reconciliation calendars, and create compliance windows measured in years, not months.

GUARD applies to Part D drugs covering oral medications typically picked up at retail pharmacies. Performance period runs January 1, 2027 through December 31, 2031, with payment years extending through December 31, 2033 for final reconciliation. The model targets sole-source drugs and sole-source biological products in specific USP therapeutic categories including analgesics, anticonvulsants, antidepressants, antineoplastics, antivirals, blood glucose regulators, cardiovascular agents, immunological agents, and respiratory agents. Approximately 25 percent of Part D enrollees selected by geographic randomization using ZCTAs.

GLOBE applies to Part B drugs covering physician-administered medications in outpatient settings. Performance period runs October 1, 2026 through September 30, 2031, with payment years through September 30, 2033. The model targets single source drugs and sole-source biological products in USP categories including anti-infective agents, antineoplastics, blood products, central nervous system agents, immunologic agents, metabolic bone disease agents, and ophthalmic agents. Only drugs with at least \$100 million in annual Medicare Part B FFS spending included. Approximately 10 percent of Medicare Part B FFS beneficiaries selected by ZCTA randomization.

Both models use the same mechanics. Manufacturers calculate a Medicare net price by taking WAC minus manufacturer rebates and discounts. CMS identifies an international benchmark using data from 19 reference countries (Australia, Austria, Belgium, Canada, Czechia, Denmark, France, Germany, Ireland, Israel, Italy, Japan, Netherlands, Norway, South Korea, Spain, Sweden, Switzerland, United Kingdom). The benchmark gets adjusted by GDP PPP factors. If Medicare net price exceeds benchmark, manufacturers owe a rebate. The rebate calculation happens quarterly with invoicing occurring no later than 6 months after the end of each calendar quarter. Manufacturers have 30 days to pay. Regular reconciliation occurs within 3 months. Additional reconciliations can happen if CMS identifies errors or manufacturer misreporting. Suggestion of error window is 10 calendar days. Administrative and judicial review explicitly precluded.

This is not about whether international reference pricing is good policy. This is about the fact that CMS just mandated quarterly benchmark calculations, HCPCS unit normalization, PPP adjustments, multi-country data integration, and deterministic reconciliation processes for a chunk of Medicare drug spend starting in less than 18 months. Every manufacturer affected needs systems. Every actuary forecasting liability needs systems. Every provider reconciling patient responsibility needs systems. That creates markets.

## **The Mechanical Reality of Both Models**

The rebate calculation requires four computational steps. First, identify the Medicare net price by subtracting manufacturer rebates from DIR and Manufacturer Discount Program discounts from WAC. Second, compute the international benchmark using either Method I (lowest GDP-adjusted country-level price from existing data sources) or Method II (volume-weighted average of manufacturer submitted net pricing data). Third, compare Medicare net price to benchmark. Fourth, if Medicare net price exceeds benchmark, calculate per unit rebate amount, multiply by total billing units during the quarter, apply any shortage reductions, and generate invoice.

For GLOBE, there's an extra step. When the rebate triggers, beneficiary coinsurance gets reduced. If the benchmark plus add-on is lower than the specified amount, coinsurance percentage drops below 20 percent using a formula that scales by the ratio. Providers must reduce what they charge patients and Medicare payment accordingly upward. That means claims processing systems need logic to detect GLOBE drugs, identify GLOBE beneficiaries, calculate adjusted coinsurance, modify remittance, and track patient collections. This happens in real time at point of payment, not quarterly during rebate reconciliation.

For GUARD, the model only applies to Part D drugs dispensed to beneficiaries in model geographic areas. That means utilization forecasting requires estimating geographic penetration not just national volume. Manufacturers can't just assume pro-rata share. They need to infer from claims patterns which percentage of their drug spend flows through the model cohort. For Part D, claims data shows beneficiary address but manufacturers don't see individual addresses. They see aggregate utilization. Forecasting model liability becomes a segmentation problem with structural uncertainty.

Both models use billing units as atomic calculation unit. For GLOBE, that's HCPCS billing units based on the HCPCS Level II code descriptor. For GUARD, that's units tied to NDC-9 level expressed in NCPDP units. Unit conversion matters because international data sources use different packaging standards, dose measurement strength conventions. A drug approved as 50mg/2mL vial in US might appear as 25mg/mL solution in Europe. Benchmark calculation requires converting all

presentations into common unit basis, applying that consistently across countries defending the math when manufacturer submits suggestion of error.

The reference country set is fixed. CMS identified 19 countries in October 2025 CIA World Factbook data for 2024. Those 19 countries remain unchanged for entire model duration even if economic indicators shift. That reduces policy volatility but creates data sourcing problems when pricing information for specific drugs becomes unavailable in specific countries during performance years. The regulations explicitly allow CMS to use the most recent available pricing data going back to certain earlier periods, but manufacturers forecasting rebates need to model scenarios where benchmark relies on stale information.

GDP PPP adjusters are static. CMS uses 2024 values from CIA World Factbook and applies them throughout model period. Currencies fluctuate. Purchasing power changes. But the adjuster stays locked. That simplifies calculation but means benchmark no longer represents real purchasing power by performance year 5. It represents purchasing power as of 2024 adjusted by nominal exchange rates where source converts to USD. This matters for Method II submissions where manufacturers report net prices in local currency and CMS applies adjuster.

## **Why These Models Create Actual Markets Instead of Consulting Theater**

Most health tech funding goes to companies solving problems that are vague enough to require selling rather than obvious enough to require buying. GUARD and GLOBE reverse that. The problems are specific. Manufacturers must calculate benchmark face penalties. Finance teams must accrue quarterly liability or fail audit. Providers must adjust coinsurance or over-collect from patients. The regs define what compliance looks like and when it's due. That makes the buying decision obvious to the people with budget authority.

The enforcement mechanism is real. GUARD and GLOBE both exclude administrative and judicial review for most decisions. Suggestion of error is discretionary and narrow. CMPs apply for failure to pay. Manufacturers can't lit

their way out. They have to operationalize. That means systems get purchased not to make internal stakeholders feel innovative but because legal and finance won't switch off on manual processes for multi-million dollar quarterly liability calculations and the audit trail is "we used a spreadsheet."

The timelines are compressed. For GLOBE, model starts October 2026. First rebate invoice February 2027. First reconciliation happens within 12 months. For GUA model starts January 2027. First rebates invoice by July 2027 for Q1 usage. Manufacturers affected by both models can't wait for requirements to stabilize. Requirements are in Federal Register now. Procurement cycles for enterprise software take 6-12 months. That means buying decisions happen in first half of 2026 for delivery later that year.

The data doesn't exist yet in usable form. Manufacturers have ASP submissions but those don't map cleanly to international analog identification. They have DIR data that's in PBM formats not benchmark calculation formats. They have pricing information from some countries but not consistent coverage across all 19 reference countries. They have costing systems but those were built for GAAP reporting not unit Medicare net price calculation by NDC-9 or HCPCS. The gap between existing systems and compliance requirements creates implementation revenue.

## **Business Model One: Benchmark Replication as Mandatory Infrastructure**

The first market is benchmark replication engines. These are not analytics. These are calculation systems that produce audit-defensible benchmark values quarterly, with version control, with provenance, with ability to replay historical calculations exactly as run, and with explanation output that legal counsel can submit during suggestion of error window without needing to translate "the model said so" into regulatory language.

For Method I benchmarks, the system needs data ingest from whatever sources (uses IQVIA MIDAS, GlobalData POLI, Eversana NAVLIN, or similar), NDC/HC to international analog mapping with strength and dosage form normalization,

presentation-level unit conversion to billing units, price aggregation by country volume weighting if available, GDP PPP adjuster application, outlier filtering by percent of US average price for GLOBE, country-level price identification, mini selection, and rounding to specified decimal places. Every step needs to be traceable. Every input needs timestamp and source citation. Every calculation needs to be replayable.

For Method II benchmarks, the system needs secure enclave for manufacturer submitted data under data agreement terms, validation logic to verify completeness and reasonableness, volume-weighting across countries, GDP PPP adjuster application, sanity checks against Method I to flag submissions that seem wrong calculation of across-country average. Then comparison of Method II to Method I to identify which is greater. All of this with audit logs showing who accessed what and what validation checks ran.

The customer is government pricing and market access finance teams at manufacturers, controllership teams responsible for quarterly accruals, and advisory firms handling CMS model compliance. The wedge is “model liability estimate for quarter close” which ties into existing financial reporting cadence. The defensibility comes from being able to explain benchmark deltas at unit level when invoices are sent and manufacturer has 10 days to submit suggestion of error or 30 days to pay.

Pricing is annual subscription based on number of covered NDC-9s or HCPCS codes plus model coverage (GUARD only, GLOBE only, or both), with implementation for portfolio mapping and system integration. The sticky upsell is managed service for rapid response during suggestion of error windows because that’s where customer pain concentrates into acute incidents with executive visibility.

The moat is not UI. The moat is regulatory knowledge embedded in calculation logic. When CMS updates guidance or issues clarifications, the system updates calculation logic automatically with version tracking. When manufacturers upgrade, they don’t just get new features, they get continued compliance with latest CMS interpretation. The switching cost is explaining to audit committee why the company fired the s

that had institutional knowledge about how CMS interprets “presentation level” hired a new vendor that has to learn from scratch.

## **Business Model Two: Rebate Accrual Systems That Survive Quarter Close**

The second market is rebate accrual and cash planning systems. Manufacturers and subledgers that estimate quarterly liability, book accrual entries, track expected invoices, post actual invoices, manage variance analysis, handle reconciliation adjustments over multi-year tail, and survive SOX controls audit.

The complexity is temporal. GUARD payment period extends through 2033 even though performance period ends 2031. GLOBE payment period extends through even though performance period ends 2031. That’s because invoicing lags quarter up to 6 months, reconciliation happens within 12 months after invoice, and CMS reopen for errors or misreporting. An accrual booked in Q1 2027 might get reopened in Q2 2028, reopened for CMS error in Q4 2028, and finalized in Q1 2029. The system has to carry open items through the tail without losing transaction history.

The second complexity is estimation uncertainty. For GUARD, manufacturers must forecast utilization within cohort. For GLOBE, manufacturers must forecast which beneficiaries receive drugs in model geographies. Both require actuarial modeling, not just trend analysis. The system needs to support scenario planning, confidence intervals, and documentation of assumptions because when actual diverges from estimate by material amount, management needs to explain why in earnings calls and audit committee meetings.

The technical architecture is event sourced ledger with rules engine. Minimum version uses quarterly events keyed by drug, model, quarter, estimated benchmark, estimated units, calculated accrual, actual benchmark when received, actual units when determined, actual invoice amount, variance, reconciliation adjustments, and final settlement. The mature version integrates with ERP via journal entry APIs, supports role-based approval workflows, provides immutable audit trail, and

generates evidence packets tying each accrual to input datasets and calculation versions.

The customer is controllable and government pricing finance. The pitch is “su quarter close without heroics” which resonates because right now most companies run these calculations in spreadsheets that break when formulas cascade errors through linked worksheets. The ROI is avoiding material weakness findings from external audit and reducing unplanned write-offs when reconciliation reveals accounts were wrong.

Pricing is annual subscription with implementation fee, potentially with shared savings component tied to variance reduction (that is, if the system’s forecasts reduce the average absolute error between estimated and actual invoices, customer shares the working capital benefit). The defensibility comes from being the system of record where every rebate artifact lives: estimates, invoices, payments, reconciliations, error corrections, shortage adjustments.

## **Business Model Three: Provider-Side Coinsurance Correctness Tools**

The third market is provider-side coinsurance compliance for GLOBE. Providers aren’t model participants but they feel the model when claims process with adjusted coinsurance and remittances show different patient responsibility than expected. Revenue cycle systems need logic to detect model drugs, confirm model beneficiary status, calculate adjusted coinsurance, validate remittance matches expectation, handle secondary payer coordination, and identify over-collections requiring refund.

The challenge is that providers don’t get advance notification of beneficiary model status. CMS updates GLOBE Eligible Beneficiary List weekly but providers don’t receive the list. They find out retrospectively when claim processes and remittance shows adjusted amount. That means the detection logic has to infer from remittance patterns and claim adjudication markers rather than checking against beneficiary roster.

The second challenge is coinsurance percentage varies by drug and quarter because depends on benchmark calculation. A drug with low benchmark has higher coinsurance reduction. A drug with high benchmark might have no adjustment. system can't hard-code percentages. It needs to learn from remittances, build lookup table of drug-quarter-coinsurance mappings, and apply those to estimate expected patient responsibility before claim processes so front desk can quote accurate amounts.

The product is claims and remittance analytics that compares actual patient responsibility to expected patient responsibility, flags discrepancies, generates workflows, provides patient statements with correct amounts, and produces audit reports showing coinsurance compliance by drug, by quarter, by location. The data sources are claim files, remittance files, and potentially GLOBE Model website information if CMS posts drug lists and benchmark amounts.

The customer is revenue cycle operations at large health systems and large specialty practices (oncology, ophthalmology, rheumatology) that administer high volumes of GLOBE drugs. The pitch is “avoid over-collecting from patients and avoid under-collecting from secondary payers” which matters because state insurance commissioners care about provider billing accuracy and patients complain when statements are wrong.

Pricing is per facility or per million in allowed charges with potential shared savings component tied to detected over-collections and prevented refunds. The defensibility comes from building high-fidelity map of remittance behaviors and adapting quickly when CMS issues operational clarifications about how coinsurance adjustment appears in claim processing.

## **Business Model Four: Unit Normalization as Boring Infrastructure**

The fourth market is HCPCS/NCPDP unit truth. Every benchmark calculation requires converting international presentation-level quantities into Medicare billable units. Every drug data vendor claims they can do this. The reality is it's a minefield

packaging changes, descriptor updates, strength differences, dosage form variations and quarter-specific mapping that breaks when manufacturers reformulate or with FDA updates HCPCS descriptors.

The product is versioned mapping service with API that returns, for any drug identifier and presentation description, the expected billing code, units per pack conversion logic, effective dates, and provenance. This is not about clinical interoperability. This is about billing semantics. It's infrastructure that makes "50mg/2mL vial maps to 25mg billing unit" into a service call not a manual lookup in NDC directory.

The technical architecture is knowledge graph or relational mapping store with temporal validity, backed by curation workflow that captures evidence for each mapping, and test harness that replays known examples and flags drift when source descriptors change. The audit value is that when benchmark calculation divides presentation-level price by billing units per presentation, the system can explain that divisor is what it is with citations to FDA documentation, manufacturer label and NDC directory entries.

The customer is benchmark replication vendors (who need reliable unit conversion dependency), manufacturers (who need conversion for internal forecasting), and potentially CMS contractors (who need conversion for invoice calculation validation). The pitch is "don't let unit conversion errors blow up your reconciliation" which matters because suggestion of error window is 10 days and arguing about unit definitions takes longer than 10 days.

Pricing is usage-based API with premium support for exceptions and portfolio onboarding. The defensibility comes from being citational. Every mapping has evidence. Every conversion has audit trail. When disputes happen, the system produces the receipts.

## **Business Model Five: Shortage Intelligence and Supply Chain Disruption**

# Tracking

The fifth market is shortage adjustment intelligence. Both GUARD and GLOBE reduce rebate amounts for drugs in shortage and for severe supply chain disrupt using formulas that count days in shortage during quarter and apply percent reductions. A single shortage event can create large financial swing for high-spe drug and large operational scramble.

The product is daily tracking of shortage list status with effective date ranges, mapping to drug concepts and billing codes, quarter day fraction calculation, pu accrual systems, and scenario planning. If drug enters shortage mid-quarter, what the expected reduction to rebate liability? This is specialized data product but it value because it ties directly to defined formulas in regulations and influences ca

The technical architecture is time series event ingestion with strong entity resolution and quarter calendar engine. The defensibility comes from clean mapping and evidence capture. When accrual changes because shortage status changed, financial teams need to document why variance occurred and what data supported the adjustment.

The customer is government pricing and controllership at manufacturers and potentially payers who need to forecast manufacturer behavior. The pitch is “know your shortage exposure before quarter close” which matters because unplanned adjustments create earnings surprises.

Pricing is annual subscription with optional scenario planning tier. The defensibility is being single source of truth for shortage status with historical tracking and forecast alerts.

## **Business Model Six: International Data Submission Management**

The sixth market is international data submission management for Method II benchmarks. Manufacturers who elect to submit net pricing data must aggregate

by country, apply volume weighting, calculate across-country averages, apply GE PPP adjusters, attest to accuracy, submit within 30 days after applicable ASP calendar quarter end for GLOBE or 180 days after performance year end for GUARD, and maintain data agreement compliance.

The challenge is manufacturers don't have systems built for this. Their pricing databases are regional or country-specific. Their rebate tracking is separate from revenue accounting. Their cost systems don't align to Medicare billing units. But submission capability requires integrating across multiple source systems, implementing currency conversion, applying regulatory calculation logic, generating attestation reports, and maintaining audit trails.

The product is data aggregation engine that pulls from manufacturer systems, normalizes to GLOBE/GUARD requirements, calculates required values, validates outputs against reasonableness checks, generates submission files, tracks submission status, and maintains history for reconciliation support. This is ETL plus compliance logic plus workflow automation.

The customer is government pricing operations and IT at manufacturers who are submitting Method II data is worth potential rebate savings. The pitch is "submit Method II without building internal capability" which matters because manufacturers are evaluating whether to elect optional submission need to factor implementation cost into ROI calculation.

Pricing is annual subscription with implementation fee for source system integration. The defensibility is regulatory knowledge embedded in calculation logic and validation rules. When CMS updates submission requirements, the system updates automatically.

## **Technical Architecture That Passes Federal Audit**

The architecture pattern that survives is replayable calculation pipeline with immutable inputs, versioned rules, deterministic outputs, and human-readable

explanation.

For benchmark calculation, this means ingest layer storing raw pricing data, HC descriptors, NDC directories, GDP PPP values, and claims-derived utilization as append-only datasets with effective dates and source metadata. Normalization layer transforms raw inputs into canonical entities (country, drug, presentation, billing unit, quarter). Calculation engine is pure functions where same inputs and same version always produce same outputs. Evidence generator outputs calculation trail that legal can read.

For Method I, the calculation engine explicitly executes the steps CMS describes: compute billing units in presentation, compute price per billing unit, filter outliers, compute country averages, apply GDP PPP adjusters, select minimum. Even if some steps only apply to certain data source types, the system represents them explicitly so it can adapt when data sources change.

For rebate calculation, the engine supports branching logic: compare benchmark-based difference to inflation-adjusted difference, use greater value, compute incremental amount, multiply by billing units, apply shortage adjustments. That branching produces subtle bugs if implemented ad hoc. It should be declarative and set or heavily unit tested code with golden datasets.

On workflow side, the system designs around reporting sequence CMS describes: There's Medicare Part B Drug Inflation Rebate Program invoice, then GLOBE/GUARD invoice for incremental amounts if applicable, then payment due within 30 days, then reconciliation reporting. The product treats each report as an object triggering tasks: validate inputs, compute expected amount, compare to received invoice, open discrepancy case, prepare suggestion of error packet with 30-day window if needed.

The system does not pretend it can appeal. The regulations explicitly preclude administrative and judicial review. The product focuses on mathematical correctness and evidence collection not legal theater.

# Where This Breaks in Practice and How to Build Around It

The big risk is not that teams can't compute benchmarks. The risk is that everyone computes slightly different benchmarks and spends money arguing about whose rounding is more spiritual. CMS specifies rounding at various points (fifth decimal place for GDP-adjusted prices in GLOBE). The more the system locks down the details and replays them exactly, the fewer disputes arise.

The second risk is cohort estimation. For GUARD, CMS randomly selects 25 percent of Part D enrollees by ZCTA. For GLOBE, CMS randomly selects 25 percent of 1 FFS beneficiaries by ZCTA. Manufacturers don't get the beneficiary list. They must infer from claims patterns which percentage of their utilization flows through each cohort. Forecasting model liability requires forecasting cohort penetration not just national volume. Good products treat that uncertainty honestly with ranges and drivers, not magic precision.

The third risk is data availability. For Method I benchmarks, CMS uses existing sources. If pricing information for GLOBE/GUARD drug isn't available in existing sources for certain countries, benchmark gets calculated with remaining countries. Manufacturers forecasting rebates need to model scenarios where benchmark relies on subset of reference countries or stale information from prior quarters.

The fourth risk is unit conversion errors. International packaging differs from US packaging. Strength varies. Dosage forms diverge. Converting everything to common billing unit basis requires judgment calls about whether presentations are truly comparable. Those judgment calls become disputes during suggestion of error. Products that document reasoning and cite sources for each conversion decision survive disputes better than products that treat unit conversion as automatic.

The fifth risk is timeline pressure. Manufacturers have 30 days to pay after receiving invoice. Suggestion of error is 10 calendar days. That's not much time to validate invoice, identify discrepancies, gather evidence, and submit formal challenge.

Products that automate invoice validation and pre-populate suggestion of error templates help manufacturers meet deadlines.

The sixth risk is regulatory drift. CMS issues guidance, clarifications, and corrections. Calculation logic needs to adapt. Products that treat regulatory content as versioned data that drives calculation rules can update automatically when guidance changes. Products that hard-code logic based on initial reading of Federal Register freeze technical debt into architecture.

## **Why Boring Infrastructure Wins When Policy Gets Specific**

The health tech funding market consistently rewards storytelling over execution because most healthcare problems are ambiguous enough that you can sell vision without proving delivery. GUARD and GLOBE reverse that dynamic. The problem is specific. The deliverable is defined. The timeline is compressed. The penalty is real.

That creates market for companies that do the boring work. Map NDCs to international analogs. Normalize presentation quantities to billing units. Apply PPP adjusters correctly. Track shortage status daily. Maintain immutable audit log. Generate calculation traces that legal can submit. Integrate with ERP. Support SaaS controls. Do it all deterministically so the same inputs produce the same outputs every time.

The winners won't be companies that build beautiful dashboards showing international price comparisons. The winners will be companies that become systems of record for benchmark calculations, invoice validation, reconciliation tracking, coinsurance adjustment. They'll be companies where switching cost is explained to auditors why the firm fired the system that had two years of calculation history and hired a new vendor starting from scratch.

The market size is every manufacturer affected by GUARD or GLOBE or both, every provider administering GLOBE drugs and handling coinsurance adjustments, every PBM managing DIR for GUARD drugs, and every advisory firm helping clients

navigate model compliance. That's not a niche. That's a forcing function where regulatory specificity creates mandatory infrastructure demand.

The funny part is the market will try to fund shiny analytical tools first because what health tech investors pattern match to. The boring part is that the real revenue will go to whoever becomes the source of truth for calculations that must happen quarterly and must be correct because the invoice is due in 30 days and the suggestion of error window is 10 days and nobody wants to explain to the CFO why we're accruing unplanned rebate liability because our unit conversion was wrong.





5 Likes

← Previous

Next

## Discussion about this post

Comments

Restacks



Write a comment...