Medicare Implant Pricing and the $25 Billion Blind Spot: Why Devices Escaped IRA-Style Negotiation, What Australia and Japan Already Publish, and What a US Device Price Rule Would Actually Require

AUG 08, 2026

🎧 Part I Podcast free on Spotify.

Thoughts on Healthcare Markets & Technology

Part I: Medicare Implant Pricing & the $25 Billion Blind Spot: Why Devices Escaped IRA-Style Negotiation, What Australia & Japan Already Publish, & What a US Device Price Rule Would Actually Require

Medicare pays for ~1 million implant procedures a year and does not know, at the product level, what it bought. No brand. No model. No manufacturer. The claim has no field for it...

Listen now

2 hours ago · Thoughts on Healthcare

🎧 Part II Podcast episode for paid subscribers only. Also available on Spotify.

Thoughts on Healthcare Markets & Technology

Part II: Medicare Implant Pricing & the $25 Billion Blind Spot: Why Devices Escaped IRA-Style Negotiation, What Australia & Japan Already

Publish, & What a US Device Price Rule Would Actually Require

Medicare pays for ~1 million implant procedures a year and does not know, at the product level, what it bought. No brand. No model. No manufacturer. The claim has no field

Listen now

2 hours ago · Thoughts on Healthcare

To listen to paid episodes in Apple or Spotify, li
settings on those platforms (instructions inside th
Podcast).

Table of Contents

1. The tweet, and the one thing in it that's
2. Where the $25 billion actually comes fro
3. Medicare doesn't buy implants, and that
4. Why the IRA walked right past devices
5. The gag clause era, and the lawsuits nob
6. Australia publishes it. Japan publishes it. France publishes it.
7. The claim form is the crime scene
8. What CMS just did in the FY2027 rule, and what it pointedly did not do
9. Bundles as the workaround, and why CJR-X is the real device pricing policy
10. What a device price rule would actually have to say
11. Four ways this goes sideways
12. What to watch

Abstract

- A public back-and-forth about hip implant costs has resurfaced a 15-year-old policy failure: nobody outside a hospital supply chain office can find out what a

US hospital pays for an implant.

- Medicare spending exposed to implantable device acquisition costs plausibly lands in the $18B to $25B range annually across ortho, spine, cardiac rhythm, structural heart, neuromodulation, and vascular.
- Devices escaped IRA-style negotiation for structural reasons, not just lobbying: no ASP reporting, no NDC-equivalent on claims, no rebate program, and payment that is bundled rather than per-unit.
- Documented price spreads on identical constructs run 78% to 83% between hospitals, with 90th vs 10th percentile ratios of roughly 2.1x for knees.
- The “no country publishes this” claim is false. Australia publishes minimum benefits for 11,000-plus device items. Japan lists roughly 200,000 items across about 1,200 functional categories and repricies them biennially off a transaction-price survey.
- The FY2027 IPPS final rule adds a UDI capture measure and nationalizes joint replacement bundles, but still leaves device identifiers off the claim.
- A workable rule needs four pieces: identifier on claim, manufacturer net price reporting, a gag clause ban, and outcome data pairing.

The tweet, and the one thing in it that’s wrong

The framing is roughly this: someone with a large audience and a functioning spreadsheet goes looking for what a hip implant actually costs, cannot find it, and says so publicly. The estimate floated is that Medicare “allocates” something like $24 billion for the implants themselves. Comparisons get made to the UK, Canada, Australia. Cue several hundred replies from supply chain directors who know the answer and cannot say it out loud without violating a contract.

The reaction from health policy people has been a little smug, which is a mistake, because the underlying observation is correct and has been correct since at least 2012. The one part that is genuinely wrong is the word “allocates.” Medicare does not allocate anything for implants. There is no line item. There is no device budget. There is no CMS office that has ever negotiated the price of an acetabular cup. Medicare pays a hospital or an ASC a bundled facility payment for a procedure, and what happens to that money inside the four walls of the facility is, from Medicare’s perspective, entirely the facility’s business.

That distinction matters enormously, because it explains why every instinct imported from the drug pricing fight fails on contact with devices. Drug pricing debates assume

a payer paying per unit for an identified product with a reported average price. Device pricing has none of those four things. Not the per-unit payment, not the identified product, not the reported average, and in most settings not even a payer who knows the product was used.

So the good news is that the question is the right question. The bad news is that answering it requires rebuilding plumbing that has been broken since the Reagan administration.

Where the $25 billion actually comes from

Nobody publishes this number, so here is the arithmetic, which anyone with a claims file and a supply contract can reconstruct within a reasonable band.

Start with lower extremity joint replacement. Traditional Medicare plus Medicare Advantage runs somewhere in the neighborhood of 800,000 to a million hip and knee replacements a year, now split messily across inpatient, hospital outpatient, and ASC after knees came off the inpatient-only list in 2018 and hips followed in 2020. Implant constructs for a primary case cluster between $4,000 and $6,000 for the big brands, cheaper for commodity and private label, considerably more for revision. Call it $4 billion to $5 billion.

Spinal fusion adds another chunk. Medicare volume runs in the low hundreds of thousands per year, and hardware for a multilevel construct lands anywhere from $8,000 to $15,000 once you count screws, rods, interbody cages, and biologics. That is comfortably $2 billion to $3 billion, and the FY2027 rule's creation of new extensive and complex spinal fusion DRGs is a tell that CMS knows the cost distribution has gotten weird.

Cardiac rhythm management is where the per-unit numbers get genuinely silly. Pacemakers run a few thousand to eight thousand. Implantable defibrillators, per the last serious federal look at this, had median hospital acquisition prices in the $16,445 to $19,007 range depending on model. Across pacemakers, ICDs, and CRT devices in the Medicare population, another $3 billion is not a stretch.

Then structural heart. Transcatheter aortic valve replacement is now the dominant aortic valve intervention in the over-75 population, and the valve itself carries a list price north of $30,000. At current US volumes, and given that this is an almost entirely Medicare-aged procedure, that alone is a $3 billion line. Left atrial appendage

occlusion adds another billion-ish. Mitral and tricuspid transcatheter therapies are climbing fast off a smaller base.

Add coronary stents, endovascular grafts, spinal cord stimulators at $20,000 to $25,000 a system, shoulder arthroplasty (the fastest-growing arthroplasty segment nobody talks about), trauma hardware, hernia mesh, and roughly four million cataract procedures a year where the intraocular lens is a small unit price times an enormous denominator.

The honest range is $18 billion to $25 billion of Medicare-exposed implantable device acquisition cost per year. So the number floating around social media is not crazy. It is in the right order of magnitude, arrived at by the wrong mechanism, which is a pretty good description of most public healthcare cost discourse.

For scale: that is larger than what Medicare spent on the first tranche of drugs selected for negotiation under the Inflation Reduction Act. The drug pricing fight got a statute, a negotiation program, a CMS office, and roughly a decade of congressional oxygen. The device pricing fight got a GAO report in 2012 and a lot of nodding.

Medicare doesn’t buy implants, and that is the entire problem

Here is the structure that makes device pricing weird, and it is worth being precise about it because most commentary gets it backwards.

The hospital buys the implant. The surgeon picks the implant. The patient receives the implant. Medicare pays the hospital a flat amount for the procedure regardless of which implant was picked. Four parties, and the only one with any price sensitivity is the one with the least control over the decision.

This is the classic physician preference item dynamic, and it has survived forty years of attempted reform because it is load-bearing for hospital-surgeon relationships. A hospital that tries to standardize to two vendors is a hospital that risks losing a high-volume arthroplasty surgeon to the joint venture ASC across the street, and with him three hundred cases a year of very profitable outpatient volume. Vendor reps are in the OR. They know the tray. They know the surgeon’s preferred offset. Try telling a spine surgeon he is switching pedicle screw systems to save the hospital $2,000 a case and see how that conversation goes.

Now run the arithmetic on what that does to margin. Medicare pays roughly $13,500 for an inpatient primary knee under the major joint DRG without major complications. Hospital outpatient lands around $12,200 in Medicare payment on a

total procedure cost near $14,100. The ASC rate for a total knee sits near $10,500 in Medicare payment. Against those numbers, a $5,000 implant is 37% of the inpatient payment and roughly 48% of the ASC payment.

Compare that to what the surgeon gets. Total knee and total hip both carry 19.11 work RVUs, which after the conversion factor and the 90-day global period works out to something like $1,100 to $1,250 for the operation and three months of follow-up. So the piece of metal and polyethylene costs about four times what the person who implants it gets paid, and the piece of metal has never once been called at 2am about a swollen calf.

That ratio is the single most important number in orthopedic economics and almost nobody outside the service line knows it.

Why the IRA walked right past devices

There is a lazy explanation for why devices got excluded from Medicare drug price negotiation, which is that the device lobby is good at its job. That is true and insufficient. The statute could not have been extended to devices without inventing four things that do not exist.

First, an average sales price. Drug manufacturers report ASP quarterly under a specific section of the Social Security Act, net of most discounts and rebates, and Part B pays ASP plus six percent off that reported number. There is no analogous reporting obligation anywhere in device regulation. The FDA regulates device safety and effectiveness and has essentially no interest in what anything costs. CMS has never asked, except in the narrow case of transitional pass-through applications where it does collect invoice cost, which is a fact that will become important later.

Second, a rebate infrastructure. The Medicaid Drug Rebate Program forced manufacturers to calculate average manufacturer price and best price, which created decades of pricing data as a byproduct of a completely different policy goal. Devices never had that. There is no device best price. There is no device AMP. There is nothing to build a negotiation formula on top of.

Third, a product identifier that appears on a claim. Drugs have had national drug codes since the early 1970s. Every dispensed drug carries one, it flows through the pharmacy claim, and it is the reason you can build a pricing dashboard on a Tuesday afternoon. Devices got unique device identifiers in a 2013 FDA rule, and the identifiers are real and populated in a public database, but they are almost entirely absent from claims. More on that shortly.

Fourth, a stable unit to negotiate. A small molecule is the same molecule for twenty years. An orthopedic implant system gets a revised tibial baseplate every eighteen months, a new porous coating, a robotic-assisted version, a cementless option, and a new brand name. Negotiating a three-year price for a product generation that will be commercially obsolete in two is not obviously coherent. This is a real objection, not just industry cover, and it argues for reference pricing or competitive bidding rather than product-by-product negotiation.

Put those together and the IRA exclusion looks less like a carve-out and more like an admission that the prerequisites were missing. Which they were. Which they still are.

The gag clause era, and the lawsuits nobody remembers

The last time the federal government looked hard at this, it found what everyone expected and then the story died.

The 2012 GAO work surveyed hospitals, group purchasing organizations, and the federal health systems on what they paid for five categories of implantable device. The results are still the best public data anyone has. One hospital reported paying about $4,500 for a primary total hip construct while another paid about $8,000 for the same construct, a 78% spread. One paid about $5,200 for a primary total knee construct against another at $9,500, an 83% spread. Identical defibrillator models showed hospital-to-hospital differences of $6,844 to $8,723. Spending on procedures involving implantable devices rose from $16.1 billion to $19.8 billion over the five years GAO examined, with orthopedics driving most of the growth.

Later work using a larger hospital sample found the 90th percentile hospital paying roughly 2.1 times the 10th percentile for knee implants and about 1.7 times for hips. The same analysis found that hospitals using a joint hospital-physician committee for vendor selection paid 17% less for knees and 23% less for hips, which is the closest thing to a free lunch in this entire domain and which almost nobody implemented.

Then there is the payer side, which is where the arbitrage lives. A widely cited analysis of commercial claims found insurers paying roughly $10,605 for knee implants that hospitals acquired for about $5,023, and $11,751 for hips acquired at about $5,620, adding up to several hundred million dollars of spread in a single commercial book over a few years. Under Medicare’s bundled facility payment, that markup does not exist as a separate line, which is either a feature or a bug depending on whether you are the hospital.

The reason this data is fifteen years old rather than fifteen months old is the confidentiality clause. Device purchase agreements routinely bar hospitals from disclosing net prices, including to benchmarking services. When intermediaries built comparative price databases from hospital-supplied data anyway, at least one large cardiac device manufacturer sued them on trade secret and tortious interference theories. The cases settled before any court ruled on the merits, which was the optimal outcome for the manufacturer: no adverse precedent, and a permanent chill on the entire price benchmarking category. Every general counsel in the space read those complaints and drew the obvious conclusion.

So the market has a well-documented 80% price dispersion on identical products, a legal regime that punishes anyone who measures it, and a payment system that gives the buyer no reason to care. That is not a market failure. That is a market functioning exactly as designed by the party with the most to gain.

Australia publishes it. Japan publishes it. France publishes it.

The claim that no country publishes device prices is the weakest link in the popular version of this argument, and it is worth demolishing because the counterexamples are also the policy roadmap.

Australia maintains what is now called the Prescribed List, a legislative instrument specifying the minimum benefit private health insurers must pay for each listed device. It covers more than 11,000 items across cardiac devices, hips, knees, intraocular lenses, human tissue products, and the accessories needed to implant them. Every benefit amount is public. Every billing code is public. The list is updated multiple times a year, with the current version dated mid-2026. It is not a survey or an estimate. It is a published price schedule you can download.

Australia also did something more interesting than publishing: it used the published list as a policy instrument. A 2022 agreement between the government and the device industry association set out a methodology for benefit reductions targeting roughly $800 million to $900 million in savings over four years, with the largest tranche landing in July 2022 and benefit cuts applied to about 51% of listed items. That is what happens when you have a published price and the political will to move it.

Japan goes further and is the more instructive model for anyone thinking about US mechanics. Japanese national health insurance lists roughly 200,000 device product items grouped into about 1,200 functional categories, each with a reimbursement price. The prices get revised biennially based on a market price survey of what

hospitals actually paid, which is the device equivalent of ASP reporting, executed as a government survey rather than a manufacturer obligation. On top of that sits the foreign average price rule, which benchmarks Japanese prices against list prices in the US, UK, Germany, France, and Australia, with outlier trimming to prevent the US number from dragging the average up. If the highest reference price is more than 2.5 times the lowest, it gets dropped entirely. What remains gets capped relative to the mean, and the final Japanese price is held at a multiple of the adjusted average, with a floor preventing any single revision from cutting more than 25%.

Read that again with an American eye. Japan has built a device pricing system whose explicit design assumption is that the US price is an outlier to be excluded from any reasonable average. That is not commentary. That is codified in the pricing formula.

France maintains a published list of reimbursable products and services with tariffs negotiated by a government pricing committee. Germany does not publish device prices directly but requires a sample of hospitals to submit detailed cost accounting data, including material costs, to the institute that calibrates the DRG weights, which means German DRG rates are built from actual implant cost data rather than charge-to-cost ratios. The UK went at it from the clinical side, with a national improvement program that published implant price variation across trusts and pushed toward standard high-performing implants, backed by a national joint registry and a public rating system that scores implants on documented survivorship.

That last piece is the one Americans keep forgetting.

The claim form is the crime scene

Everything above is downstream of a single mundane fact: the US institutional claim has nowhere to put a device identifier.

The unique device identifier system has existed since 2013. Implantable devices carry them. They are registered in a public FDA database with manufacturer, model, brand, and attributes. Certified EHRs have been required to maintain an implantable device list keyed to UDI for a decade. The data exists, it is standardized, and it is sitting in the medical record.

It is not on the claim. The institutional claim format adopted under HIPAA has no field for it. Revenue codes for implants exist, but they are cost center buckets, not product identifiers. Outpatient device category codes exist, but they describe categories, not products. The billing committee that maintains the paper claim form

approved adding the device identifier years ago, but the electronic transaction standard that actually carries the data has no home for it, and updating the adopted standard is a multi-year regulatory process that nobody has been willing to spend political capital on.

The consequence is that CMS pays for roughly a million implant procedures a year and does not know, at the product level, what it bought. Not the brand, not the model, not the manufacturer. When a device is recalled, and they are recalled with some regularity, the agency cannot identify affected beneficiaries from claims. It has to go hospital by hospital. Two of the largest orthopedic recalls of the last two decades, one involving a metal-on-metal hip system and another involving polyethylene inserts that degraded from packaging failures, together affected hundreds of thousands of implanted patients, and the identification process in both cases was a manual scramble.

The safety argument for device identifiers on claims is stronger than the pricing argument and has been made repeatedly by former agency heads from both FDA and CMS. The pricing argument rides along behind it. You cannot construct an average sales price for a product you cannot count, and you cannot count what does not appear in any administrative dataset.

What CMS just did in the FY2027 rule, and what it pointedly did not do

Which brings us to the last few weeks, because the inpatient final rule that dropped in early August contains the most movement on this in years, and also a very deliberate stopping point.

CMS finalized a new measure in the hospital interoperability program requiring hospitals to capture the complete unique device identifier for implantable devices in the certified EHR, landing in the public health and clinical data exchange objective for the 2027 reporting period. It is attestation-based at launch, meaning hospitals answer yes or no rather than reporting a numerator and denominator, and a failure to attest costs them the payment update. The agency simultaneously issued a request for information on moving to a performance-based version and on what else UDI capture could support.

Read the placement carefully. The measure sits in the interoperability program, tied to EHR capture and public health data exchange. The stated rationale is device surveillance and safety. It does not touch the claim. A hospital can fully comply, capture every UDI perfectly in its EHR, and CMS still learns nothing about which

device was implanted in any given Medicare beneficiary, because the data never travels to the agency.

That is not an oversight. It is the maximum the agency could do without opening a transaction standard rulemaking and without picking a fight with the device industry over what the identifier would eventually be used for. There has been reporting that CMS is contemplating a proposed rule that would put device identifiers on claim forms. Contemplating. The RFI is the tell that the agency is building a record for that fight.

The same rule also proposed eliminating the alternate pathway that let breakthrough-designated devices qualify for new technology add-on payments and outpatient pass-through status without meeting the standard substantial-clinical-improvement test. That is a quiet but real tightening. Pass-through and add-on payments are the only two places in Medicare where device cost is visible as a distinct payment amount, and CMS just made them harder to get.

Bundles as the workaround, and why CJR-X is the real device pricing policy

If you cannot regulate the price, regulate the buyer's incentive. That has been CMS device strategy for a decade, and it just went nationwide.

The final rule confirmed the Comprehensive Care for Joint Replacement Expanded model, mandatory for essentially every hospital paid under the inpatient and outpatient prospective payment systems, covering hip, knee, and now ankle replacements across inpatient and hospital outpatient settings, with 90-day episodes. Start date moved to January 2028, three months later than proposed. Exclusions are narrow: hospitals already in the Transforming Episode Accountability Model, Maryland, critical access hospitals, rural emergency hospitals, and a handful of demonstration sites.

That means two mandatory episode models running simultaneously. TEAM launched in January 2026 across roughly 700-plus hospitals in about 188 statistical areas, covering five surgical episode types on 30-day windows, running through the end of 2030, at which point those hospitals roll into the expanded joint replacement model. The predecessor model produced something like $112.7 million in net savings over two performance years while holding quality steady, which is the evidence base CMS cited for going national.

Here is why this is device pricing policy in disguise. In an episode model, the implant is no longer a cost the hospital passes through to a fixed DRG. It is a controllable input inside a spend target the hospital gets measured against. The first bundled joint replacement model taught hospitals exactly this lesson, and the industry response was immediate: capitated implant pricing, matrix pricing tied to volume commitments, single-vendor agreements, and the rise of value implant vendors selling perfectly adequate primary constructs at $2,000 to $3,000 against incumbents at $5,000 to $6,000.

Make that mandatory for 2,500-plus hospitals and the aggregate purchasing pressure is larger than anything a negotiation statute would have produced in the same timeframe. Every hospital in the country simultaneously acquires a reason to care about implant cost, and the four firms that dominate large joint reconstruction acquire 2,500 counterparties who have all just read the same rule.

The catch is that bundles produce private savings, not public information. A hospital that negotiates its knee construct from $5,500 to $3,200 keeps the difference until the benchmark rebases. Nothing about that transaction becomes visible to CMS, to researchers, or to the next hospital. The dispersion problem does not get solved. It gets redistributed toward whoever has the best supply chain team, which in practice means the large systems that already had the best prices.

What a device price rule would actually have to say

Strip away the advocacy and a workable rule has four components. Three of them are boring and one of them is the whole fight.

The first is identification. Device identifiers have to appear on institutional claims for implantable devices, which requires CMS to move on the adopted transaction standard and give hospitals a real implementation window, probably three years, with the identifier populated automatically from the EHR implantable device list rather than hand-keyed by a coder. Without this, nothing else in the list is possible. With it, every other element becomes an analytics problem rather than a policy problem.

The second is price reporting. Manufacturers report average net selling price by device identifier on a quarterly lag, net of rebates, discounts, purchasing organization administrative fees, and any other consideration. This is not novel. It is the drug ASP mechanism with a different identifier, and the statutory language could be adapted in an afternoon. Alternatively, follow Japan and run it as a periodic transaction price survey of purchasers, which shifts the reporting burden to hospitals but captures the

actual paid price rather than the manufacturer's calculation of it. The invoice cost collection that already happens for pass-through applications proves the agency can handle the data.

The third is the gag clause ban. Federal law already prohibits contractual terms that prevent group health plans from accessing their own claims data, enacted as part of a broader transparency package a few years back. Extending the same principle to device purchase agreements is a one-paragraph provision that would immediately reopen the benchmarking market that litigation shut down. This is the cheapest, highest-yield item on the list and it requires no new infrastructure whatsoever.

The fourth is outcome pairing, and this is the one that separates a serious proposal from a press release. Publishing prices without survivorship data creates pressure toward the cheapest construct, and orthopedics has a long, expensive history of cheap-and-novel implants that failed at five years and generated billion-dollar settlements. The UK solved this by pairing its price transparency push with a national registry and a public rating system that grades implants on documented years of survivorship evidence. The US has a joint replacement registry with millions of procedures in it, but participation is voluntary, it does not link to price, and it does not link to claims. Device identifiers on claims would connect all three datasets automatically, which is the strongest argument for the first item on this list and the one that should be leading the pitch.

Four ways this goes sideways

Transparency in concentrated markets is not automatically pro-competitive, and anyone selling it as a sure thing is not being careful. The canonical counterexample is a Danish antitrust experiment in ready-mix concrete, where the regulator published firm-level transaction prices in a concentrated market and prices converged upward rather than downward, because publication gave each firm perfect visibility into whether rivals were discounting. Large joint reconstruction is dominated by four firms. Publishing counterparty-identified transaction prices in that structure is a plausible way to make the low-price hospitals pay more rather than making the high-price hospitals pay less. The mitigation is to publish distributions with a lag rather than named transactions, which preserves the benchmarking value while removing the enforcement mechanism for tacit coordination.

The second risk is that price becomes the only visible attribute. See above regarding cheap implants and five-year revision curves. A revision costs Medicare two to three times a primary and costs the patient a year.

The third is the surgeon relationship. Every hospital that has tried aggressive implant standardization has discovered that the surgeon is mobile and the hospital is not. Mandatory bundles help here by giving the hospital a regulatory reason rather than a margin reason, but the ASC alternative is a real and growing pressure valve, and the ambulatory surgery industry has already told CMS in writing that pulling ASCs into episode accountability would slow the migration of procedures out of higher-cost settings. They are not entirely wrong.

The fourth is innovation, which is the argument the device industry will lead with and which deserves a fair hearing even from skeptics. The US is where new device categories get paid for first, and the margin from US pricing funds the iteration. Japan's own experience is instructive here in an uncomfortable direction: as its foreign average price rule tightened and the innovation premium on new device categories compressed from roughly 10% to about 3% over a decade, the ratio of Japanese prices to the foreign average fell below parity, and the debate there is now whether the country is systematically underpaying for new technology. Whatever the US builds should have a real premium pathway for genuine clinical improvement, which makes the simultaneous tightening of the breakthrough device pathway in the FY2027 rule look a little like closing the innovation door while pretending to shop for a lock.

What to watch

Three things over the next eighteen months.

Whether CMS actually proposes device identifiers on claims. The RFI in the FY2027 rule is the strongest signal in a decade, and the safety framing gives the agency political cover that the pricing framing never would. If a proposed rule appears, the comment file will be the most interesting document in health policy that year, and the device industry's comments will tell you precisely how much the current opacity is worth to them.

Whether the mandatory bundle rollout produces visible implant price compression. Hospitals have seventeen months before the expanded joint replacement model starts. Supply chain teams are already renegotiating. The public evidence will show up first in device manufacturer earnings calls as pricing pressure commentary in large joint reconstruction, and then in the DRG relative weight recalibrations two or three years later as the cost report data catches up.

Whether the transparency bills moving through Congress pick up a device provision. Several hospital and insurer transparency bills advanced through committee this year, all of them extending the existing machine-readable file framework for services and

negotiated rates. None of them touch device acquisition cost. Adding a gag clause ban would cost nothing, score as savings, and be nearly impossible to oppose publicly. Somebody should hand that amendment to a staffer.

The larger point is that the person asking what a hip implant costs was not being naive. He was pointing at the last major category of US healthcare spending with no price data at all, in a system that has spent fifteen years building elaborate transparency machinery for hospital charges and negotiated rates while leaving 40% of the cost of a joint replacement completely dark. The drug pricing debate got a statute. The device pricing debate has a request for information and a checkbox in an interoperability program.

That is progress, technically. So is a glacier.

15 Likes • 1 Restack

← Previous

Discussion about this post

Comments Restacks

Write a comment...