Where Healthcare Cost Variance Actually Comes From

Physician and Patient Preference as Causal Drivers of Efficiency Variance

An episode-level analysis of 8.2 million commercial insurance episodes across four Texas metropolitan areas, integrated with five decades of peer-reviewed research on the drivers of healthcare cost variance.

8.2 M	**72,776**	**131**	**4**
episodes analyzed	provider care groups	episode types	Texas metros

Authors: Stephen F. Wiggins, MBA; Erica Smith PhD, MBA, MS, RN, CHDA; Emil Rusev, PhD; Earl Steinberg, MD, MPP. Oxbridge Health Analytics Team.

Audience: Employers, benefit advisors, brokers, reinsurers, health plans, TPAs, physicians, hospital executives, and healthcare industry professionals.

May 2026

Table of Contents

Executive Summary

Every employer, insurer, and government agency paying for healthcare is paying two bills at once. The first is for the price of each service. The second, larger bill is for how much care actually gets delivered. Almost nobody buying health insurance today can see those two bills separately.

This study separates them.

Using 8.2 million commercial insurance episodes across four Texas metros, grouped into 131 clinically defined episode types and attributed to 72,776 individual provider care groups, the analysis breaks total cost into a price component (what providers charge per unit of service) and an efficiency component (how many services are actually delivered). Patient status is normalized for age, sex, zip code, complications, and clinical classification. The results overturn a widely held assumption about what really drives healthcare costs.

Five Findings

1. Variance is enormous. For common conditions, the risk-adjusted cost of a single episode can differ 20x to 40x between the lowest-cost and highest-cost provider in the same city. Cataract surgery, a relatively straightforward outpatient procedure, varies roughly 9x.

2. Efficiency drives the variance. Differences in how efficiently providers deliver care account for 88% of total cost variance when episodes are weighted equally. Unit prices account for 12%. When weighted by actual dollars spent, efficiency accounts for over 99%.

3. Conditions dominate spending. Condition episodes (diabetes, depression, low back pain, asthma) represent roughly four of every five dollars spent. For conditions, essentially 100% of cost variance is driven by provider efficiency.

4. Procedures tell the opposite story. For procedural episodes (knee replacement, cataract surgery, coronary bypass), unit price drives 89% of the variance. But procedures represent only about 17% of total spending.

5. Physician and patient preference explain the efficiency variance. Physician practice style accounts for 30–35% of geographic spending variation. Patient preferences explain approximately 5%. Much of what appears as patient demand is actually physician influence operating through information asymmetry.

> *Across 8.2 million episodes in Texas, roughly 88% of the variance in total cost between providers is driven by how efficiently care is delivered, not by the prices providers are paid. For condition-based spending, that figure is effectively 100%. Physician practice patterns are the primary reason this efficiency variance exists.*

Why This Study Matters

Healthcare costs have grown faster than wages, faster than corporate revenue, and faster than general inflation for most of the last three decades. For a typical mid-sized employer, healthcare is now the second-largest expense after payroll, and the one they have the least ability to manage.

Every major buyer has tried the same levers: negotiating lower prices, narrowing networks, shifting cost to employees, introducing transparency tools. These strategies have slowed growth in places. They have not solved the problem.

One reason: they have all targeted the wrong bill.

The Institute of Medicine estimated in 2009 that approximately 30% of U.S. health spending ($750 billion) was wasted on unnecessary services, excessive administration, fraud, and other problems. A 2019 JAMA systematic review updated those figures to $760–$935 billion annually, roughly 25% of total spending. Of that, $75.7–$101.2 billion was attributed specifically to overtreatment or low-value care.

Those waste estimates did not answer who was responsible for the overtreatment or why it persisted. This study, combined with the research literature summarized below, provides that answer.

What episodes reveal that other approaches cannot

Price transparency of the kind required by federal rule is useful but incomplete. Knowing that two MRI facilities charge different rates does not tell you which physician's diabetes patients end up with half as many emergency department visits, or which orthopedic practice's total joint replacement episodes produce the most complications and the longest stays.

Episodes of care capture the full clinical journey for a single clinical problem: the physician, the facility, the testing, the drugs, the follow-up. Measured at the episode level, the variance becomes visible and actionable.

The Data Set

The study is built on a large-scale commercial medical claims data set, analyzed at the episode level.

Geography. Four Texas metropolitan areas: Dallas–Fort Worth, Houston, San Antonio, and Austin. Restricting to a single state removes differences in regulation and insurance mandates; restricting to metros removes rural-vs-urban distortions.

Claims base. Commercial medical claims representing approximately 40% of the commercial insurance market in the four metros. The sample spans condition episodes (managing an ongoing clinical problem such as heart failure or depression) and procedure episodes (a defined intervention such as knee replacement or cataract surgery). Together these represent approximately 43% of total healthcare costs for employers. Episodes beginning with an emergency room visit are excluded. Oncology episodes are excluded due to the complexity of risk adjustment.

Care groups. 72,776 provider care groups. A care group is anchored by a physician and linked to the specific facility where care was delivered. The same physician practicing at an ambulatory surgery center and a hospital is measured as two separate care groups.

Volume. 8.2 million episodes across 131 episode types covering cardiology, orthopedics, gastroenterology, behavioral health, women's health, dermatology, ENT, urology, endocrinology, and other specialties.

Risk adjustment. Each episode is risk-adjusted for age, sex, ZIP code, clinical classifications, and episode-relevant complications.

Credibility screening. Care groups below a minimum episode volume were excluded.

How the decomposition works

For every care group treating a given episode type, the study calculates two indices relative to the market average. A price index captures whether that care group's unit prices are above or below market. An efficiency index captures whether that care group delivers more or fewer services than the market average for the same risk-adjusted patient. Because total episode cost is the product of the price paid and the quantity of services delivered, the two indices multiplied together describe the full episode cost relative to market. The study decomposes total variance into the share attributable to each.

Finding 1: The Variance Is Enormous

Within the same metro area, for the same clinical problem, and after risk adjustment, the difference between the cheapest and most expensive care group for a single episode is routinely in the tens of times.

Episode	Metro	Low Price	High Price	Ratio	Care Groups
Cataract Surgery	Austin/SA	$1,311	$13,314	9.1x	169
Colonoscopy	Houston	$640	$8,006	12.5x	168
Knee Replacement	Austin/SA	$21,269	$85,052	4.0x	46
Hip Replacement	Dallas-FW	$37,043	$125,146	3.4x	65
CABG (cardiac bypass)	Dallas-FW	$106,570	$266,346	2.5x	7
Depression & Anxiety	Austin/SA	$1,522	$44,355	28.1x	1,180
Depression & Anxiety	Dallas-FW	$950	$26,492	26.8x	2,212
Acute ENT Conditions	Dallas-FW	$547	$22,882	31.7x	2,863
Contact Dermatitis	Houston	$646	$26,064	39.3x	812

A patient in Houston with the same insurance and the same clinical presentation for dermatitis can end up in an episode costing $646 or $26,064, depending on which provider they see. Joint replacements vary 3–4x across providers in the same market. Chronic and acute conditions routinely show 20–40x variation. None of this shows up on an explanation of benefits, a provider directory, or a price transparency file. It is visible only after episodes are assembled and risk-adjusted.

Finding 2: Efficiency, Not Price, Drives Most of the Variance

The central analytical exercise of this study is separating the two drivers of total episode cost.

Across all 131 episodes (equal weight)

Driver	Share of Variance	Meaning
Efficiency (volume of care delivered)	88%	Practice pattern differences between providers
Unit price (fee schedule)	12%	Contracted rate differences between providers

This result reframes the cost problem. The levers most buyers have pulled for thirty years address, at most, the 12% bucket.

Weighted by actual dollars spent

Driver	Dollar-Weighted Share	Translation
Efficiency	99.2%	Virtually all economically material variance
Unit price	0.8%	Negligible when measured in dollars

Measured against the dollars that flow through the commercial healthcare system, the fight over unit prices is a fight over roughly 1% of total spending variance.

Finding 3: Conditions and Procedures Tell Opposite Stories

Conditions: efficiency drives everything

Condition episodes	Share of Variance	Dollar Impact
Efficiency (utilization)	~100%	$5.4 billion above market
Unit price	~0% (mild offset)	($112) million offset

For conditions, the prices providers charge are roughly in line across care groups. What varies is how much care they deliver. Two primary care practices managing the same diabetic population can produce episodes differing in total cost by 10x or more, while billing at roughly market rates for each individual service. The difference is the count, mix, and intensity of services ordered.

Procedures: unit price dominates

Procedure episodes	Share of Variance	Dollar Impact
Unit price	89%	$176 million above market
Efficiency	11%	($103) million offset

Surgeries and defined interventions are more standardized in how they are delivered. The dominant variance is what providers charge per case. This is why narrow networks and centers-of-excellence programs work for procedures.

Most existing cost-management strategies were built with the procedural episode in mind: reference-based pricing, bundled payments, centers of excellence, travel-surgery programs. They work on the 17% of spending that comes from procedures. They leave largely untouched the 83% that comes from conditions, where variance is driven almost entirely by how care is delivered.

Finding 4: Why Efficiency Varies — The Role of Physician Preference

The Oxbridge study quantifies efficiency variance. The next question: why does it vary so much between providers treating the same condition in the same city?

The foundational evidence

The Dartmouth Atlas of Health Care (1993–2019) documented systematic geographic variation in medical resource use across more than 300 Hospital Referral Regions. Supply-sensitive care—services whose use varies with local resource availability rather than clinical evidence—accounted for well over half of Medicare spending and was marked by systematic overuse in high-capacity areas. Supply-sensitive spending was 66% higher in the highest-cost regions compared to the lowest-cost regions. Effective and preference-sensitive care spending was essentially flat.

These findings launched decades of research aimed at isolating the sources of variation: physicians, patients, or the system.

Physician practice style: the dominant driver

The most rigorous studies use natural experiments: physician and patient migration across regions, physician exits forcing patient reassignment, within-institution comparisons holding environment constant. Their findings are consistent.

Badinski, Finkelstein, Gentzkow, Hull, and Williams (2023, NBER) exploited both patient and physician migration to decompose regional variation. Physicians account for approximately 30–35% of geographic variation in utilization, conservatively three times as important as non-physician supply-side factors.

Cutler, Skinner, Stern, and Wennberg (2019) linked clinical vignettes to Medicare expenditure data. Physician beliefs unsupported by evidence drive 35% of end-of-life spending variation and 12% of all-enrollee spending variation. In some specifications, knowing how a small sample of area physicians would treat hypothetical patients explains over 50% of end-of-life spending variation.

Finkelstein and colleagues (2020) tracked patients who switched PCPs due to provider exit. PCP practice style explains 41–89% of primary care spending variation. When patients move to a PCP with $1 higher total spending per patient, their own spending increases by $0.48 in the first year.

Van Parys (2016) used quasi-random assignment of patients to ED physicians within a single emergency department. Physicians at the 75th percentile of spending spend 20% more than those at the 25th percentile for comparable patients, controlling for case mix and hospital effects.

Molitor (2018) tracked cardiologists who migrated across regions. 60–80% of physician behavior is shaped by the local environment; 20–40% is an intrinsic practice style that travels with the physician.

Components of physician preference

Component	Definition	Estimated Magnitude
Clinical beliefs	Beliefs about treatment efficacy not supported by evidence	35% of end-of-life spending variation (Cutler et al.)
Training and practice origin	Medical school culture, residency training embedded in practice style	20–40% of physician variance survives migration (Molitor)
Organizational culture	Peer influence and group norms within practice groups or hospitals	Elasticity of ~0.27–0.30 (Doyle)
Supply availability	Physician behavior responsive to local resource availability	60–80% of behavior explained by environment (Molitor)
Financial incentives	Fee-for-service inducement to increase volume	Magnitude varies by specialty
Patient-physician concordance	Trust dynamics modulating ordering intensity	3% increase in SDM reduces expenditures ~10% (Brown & Hurley)

Financial incentives

Under fee-for-service payment, physicians are paid more when they deliver more services. This creates a structural incentive for higher utilization that operates alongside, and sometimes reinforces, clinical beliefs about appropriate care intensity.

In maternity care, Johnson and Rehavi (NBER) found that obstetricians performed significantly more C-sections on non-physician patients than on physician patients for low-risk deliveries. A $100 increase in the fee differential between C-section and vaginal delivery was associated with a 3.4% increase in the primary C-section rate. A 2021 JAMA Network Open study found that women delivering at hospitals with higher profits per cesarean had a higher probability of undergoing C-section.

In spine surgery, the Lown Institute's 2025 analysis of Medicare claims identified more than 200,000 unnecessary back surgeries over three years, costing approximately $2 billion. Physicians who performed unnecessary procedures collectively received $64 million from device and drug companies. Sixty percent of overuse came from just 10% of worst-offender physicians.

In orthopedics, physician preference items (implants and devices) are among the largest cost drivers in joint replacement episodes. Implant costs for total knee replacement range from $1,797 to $12,093

across surgeons within the same institution, with physician loyalty to vendor relationships explaining most of the within-hospital variance.

Finding 5: Patient Preference — Smaller but Real

Physician preferences dominate patient preferences as a cost driver. But patient preference is not zero, and it operates through a distinct mechanism with real policy implications.

Baker, Bundorf, and Kessler (2014, Health Affairs) directly measured patient preferences using a survey-based instrument linked to spending at the Hospital Referral Region level. Patient preferences explain approximately 5% of total Medicare spending variation.

Patient demand-side factors (health status plus preferences combined) explain 40–50% of geographic variation when isolated by patient migration studies (Finkelstein, Gentzkow, and Williams 2016). The preference-specific share within this demand component is much smaller than the health-status share.

When patients receive full information through structured decision aids, a substantial share reverse their stated preference for surgery. For benign prostatic hyperplasia, only 14% of fully informed patients preferred surgery. This suggests that much of observed surgical variation attributed to "patient demand" is actually physician influence operating through information asymmetry.

Why physician preferences dominate

Patients depend on physician framing and recommendations. When preferences are elicited with full information, patients generally prefer less care, not more. Physicians act as agents for patients but can influence demand through diagnostic labeling, referral intensity, and treatment framing. Physicians trained or practicing in high-intensity areas adopt local norms, so physician preference is partly a collective construct patients do not independently create. In preference-sensitive care categories, local surgical capacity predicts utilization rates independently of patient demand.

Shared decision-making

Brown and Hurley (2023) published the first national study demonstrating a causal effect between shared decision-making (SDM) and healthcare expenditures. Analyzing surveys of more than 60,000 patients from 2003 to 2017, they found that a 3% increase in SDM is associated with approximately a 10% decrease in expenditures. This impact doubled for Latinx patients seen by Latinx physicians and tripled for Black patients seen by Black physicians, reflecting the role of trust and communication quality in treatment efficiency.

The implication is direct: a substantial share of current spending reflects patient choices made without full information. When that information is corrected, costs decline. Not because patients are denied care, but because fully informed patients choose less care than physicians would otherwise recommend.

> *Physician preferences dominate patient preferences as a driver of cost variance by approximately 6 to 1 (30–35% vs. 5% of geographic variation). SDM programs can reduce expenditures by 10% for every 3% increase in SDM, making correction of the information asymmetry one of the highest-return interventions available.*

Preference-Driven Variance in Practice: Episode-Level Evidence

Five episodes from the Oxbridge portfolio demonstrate the range of preference-related cost drivers.

Lumbar spine surgery

Low back pain is the single most clear-cut case. Lumbar spinal fusion is among the most expensive and fastest-growing surgical procedures in the U.S., with hospital costs reaching $12 billion in 2014. Medicare spending on inpatient back surgery more than doubled from 1997 to 2005; lumbar fusion specifically increased 500%.

A 2023 JAMA Network Open study found that in 121,745 Medicare patients undergoing elective surgery for lumbar stenosis with degenerative spondylolisthesis, decompression-alone rates fell from 32.6% in 2016 to 9.6% in 2019. Fusion rates rose from 67.4% to 90.4%. This happened despite two landmark NEJM randomized trials in 2016 finding no benefit of fusion over decompression alone. Academic surgeons kept costs 55% lower than private practice surgeons for the same conditions.

Knee and hip replacement

Implant cost is the dominant preference item. In a 29-hospital study, total knee arthroplasty costs varied 2:1 between the 90th and 10th percentile despite equal outcomes. Implant costs ranged from $1,797 to $12,093 across surgeons within the same institution. Studies implementing implant price capitation programs found cost decreases of 20–50% without adverse effects on outcomes.

Coronary artery disease and PCI

For stable coronary artery disease, the choice between percutaneous coronary intervention (PCI) and optimal medical therapy (OMT) is one of the most consequential preference decisions. Cost-effectiveness analyses show OMT has the lowest lifetime costs at approximately $22,952 versus $25,081 for bare metal stent PCI and $25,536 for drug-eluting stent PCI. The preferentially chosen interventional option is simultaneously the most expensive and not clinically superior for stable disease. Per-episode cost impact: $2,000–$12,000 for PCI over OMT.

Autoimmune conditions and biologic prescribing

Biologic drugs exceed $25,000 per patient-year, and the most expensive therapies exceed $80,000. The timing of initiation, agent selection, dosing intensity, and switching behavior are almost entirely physician- and patient-driven.

Hysterectomy and uterine fibroids

The choice of surgical approach (minimally invasive vs. open), concurrent procedures, and the decision to recommend surgery versus conservative management produce 50–68% of procedure rate variation unexplained by patient factors. Hospital length of stay adds $3,000–$6,000 per case for open versus minimally invasive approaches.

Summary by episode category

Category	Primary Preference Channel	Preference-Driven Variance
MSK Surgical	Implant selection + surgical approach	30–60% of episode cost variation
MSK Condition (Low Back)	Surgery vs. conservative + injection overuse	40–65% of episode cost variation
Cardiac	PCI vs. OMT + device selection	25–40% of episode cost variation
Autoimmune/Dermatologic	Biologic timing + agent selection	30–70% of drug cost variation
Women's Health Surgical	Surgical approach	50–68% of procedure rate variation
Maternity	C-section financial incentives	10–30% of C-section rate variation
Chronic Conditions (HF, COPD)	Post-acute routing + readmission mgmt	25–50% of 90-day episode variation

A Quantitative Framework for Healthcare Cost Variance

Variance Component	Share	Basis
Efficiency variance	88% of episode-level variance	Oxbridge Texas study (8.2M episodes)
Physician practice style/beliefs	30–35% of geographic utilization variance	Badinski et al. (2023); Cutler et al. (2019)
Other supply-side factors	~20% of geographic variance	Badinski et al. (2023)
Patient preferences (isolated)	~5% of spending variance	Baker et al. (2014)
Patient health/demographics	~12% of spending variance	Baker et al. (2014)
Price variance	12% of episode-level variance	Oxbridge Texas study (8.2M episodes)

The 88/12 split comes from the Oxbridge episode-level analysis using a single national insurer's fee schedule. Under multi-payer assumptions, the price share may rise toward 20–25%, but efficiency accounts for the dominant share under any scenario, and for conditions it remains near 100%.

The physician and patient preference percentages derive from Medicare-based geographic variation studies. These are a different unit of analysis than episode-level decomposition, but they are complementary: the Oxbridge study shows that 88% of episode cost variance is efficiency-driven, and the broader literature shows that physician preference is the primary reason efficiency varies.

Physician preferences dominate patient preferences approximately 6 to 1. Higher cost does not produce better outcomes, consistent with decades of prior research.

Strategic Options for Employers

The findings translate into concrete benefit design choices, arranged from least to most ambitious.

Option 1: Incentive-only program

The employer sets a benchmark (typically the mean or median care group price for each episode) and shares a portion of savings when employees select providers below the benchmark. This can

substantially reduce or eliminate out-of-pocket cost for many episodes. Think of it as a bell curve for each episode, with the incentive encouraging employees to select from the lower-cost portion. Savings are not contractually guaranteed and depend on employee uptake. Disruption is low; the overlay layers on top of the existing plan without a network change.

Option 2: Incentives plus guaranteed-price episodes

An upgrade from the incentive-only model. The employer and the employee know in advance what the financial outcome will be for each covered episode. The care group's price is locked in up front through a guaranteed-price arrangement, eliminating the surprise bill. Employees use a portal or mobile app to shop for guaranteed-price care groups. Savings are more predictable because guaranteed prices remove cost overrun risk.

Option 3: Allowances plus guaranteed-price episodes

The most fully developed form. An episode allowance cap is built into the benefit plan design, typically set at the 70th, 80th, or 90th percentile of care group prices in a market. Costs above the allowance are either not covered, or employers can limit additional out-of-pocket for employees who select over-allowance care groups. This design is actuarially scorable and reinsurable.

Deployment paths

Employers can offer Episode Benefit Plans as an option alongside existing plans, as a complete replacement (quoted at equivalent actuarial value for apples-to-apples comparison with PPO/HMO designs), or as an incentive-only overlay bolted onto an existing program with the cooperation of the current plan administrator.

Strategic Options for Health Plans and TPAs

The variance within the networks a plan already owns or leases is substantially larger than the variance between networks. That variance becomes the basis for a new set of product and cost-containment strategies.

Provider market consolidation makes true price competition between organizations increasingly difficult, but the data show that significant cost variance exists between care groups within the same network, and even within the same hospital or system. Fee schedule negotiation targets the smaller bucket. The larger opportunity is to surface and steer around the efficiency variance inside existing contracts.

A white-labeled episode product can deploy in approximately 120 days, with the plan retaining brand, network, contracts, and financial control. The partner provides software, analytics, and guaranteed-price infrastructure. Independently, episode-based bundled payment contracting uses episodes as the unit for payment arrangements with specialists and health systems across the plan's full product range.

Episode Overlays Work With Any Plan Design

Episode-based incentives, guaranteed-price episodes, and episode allowances can be layered on top of any underlying benefit structure: PPO, HMO, POS, reference-based pricing, or traditional indemnity. The

episode overlay addresses a dimension of cost the underlying plan does not reach. A PPO negotiates unit prices. An HMO manages utilization through gatekeeping. A reference-based pricing plan caps what it pays per service. None of them measure or manage the total cost of a complete episode of care.

The overlay does not replace the underlying plan's network, formulary, or cost-sharing structure. It adds transparency and accountability at the episode level, creating a measurable cost improvement regardless of the plan architecture underneath.

A Note on Price Data and Sensitivity Analysis

This study used the contracted payment rates of a single national health insurer, creating more homogeneity in the price component than a multi-payer blend would produce. The 88/12 efficiency-to-price split may modestly understate price variance in a multi-payer environment.

When comparing publicly disclosed fee schedules of national insurers in the Texas markets, there were modest differences in contracted price levels. Narrow networks and those with terms more favorable to providers can have meaningfully lower price levels. There are no episodes where the price paid would overcome the influence of provider efficiency on total cost variance.

Sensitivity analyses using alternative fee schedule assumptions confirm: the core finding holds. The price share may shift from 12% toward 20–25% under aggressive multi-payer assumptions, but efficiency consistently accounts for the dominant share. For condition-based spending, the efficiency share remains near 100% under any fee schedule assumption.

Conclusion

The conventional approach to healthcare cost management has focused on price. These strategies address, at best, the 12% of cost variance attributable to what providers charge per unit of service.

This study shows that the dominant driver is efficiency: how much care is delivered, by whom, in what setting, and through what combination of services. Efficiency accounts for 88% of total episode-level cost variance across 8.2 million commercial insurance episodes. For condition-based spending, it is effectively 100%.

Five decades of peer-reviewed research establish that physician practice style and beliefs are the primary reason efficiency variance exists. Physicians account for 30–35% of geographic utilization variance. Patient preferences contribute approximately 5% independently, but much of what appears as patient demand is physician influence operating through information asymmetry. SDM programs that correct this asymmetry reduce expenditures by approximately 10% for every 3% increase in SDM.

The episode of care is the unit of measurement that makes both price and efficiency variance visible, attributable, and actionable at the individual care-group level. Episode-based benefit designs address the larger portion of cost variance that conventional approaches leave untouched, and they work with any underlying plan design.

Methodology and Limitations

Study design

Commercial medical claims from four Texas metros: Dallas–Fort Worth, Houston, San Antonio, and Austin. 131 episode types spanning cardiology, orthopedics, gastroenterology, behavioral health, women's health, dermatology, ENT, urology, endocrinology, and other specialties. Each episode risk-adjusted for age, sex, ZIP code, clinical classifications, and episode-relevant complications. A care group is the combination of a physician (identified by NPI) and the facility where care was delivered. Care groups below a credibility threshold were excluded.

Analytical method

The price and efficiency indices are multiplicative, not additive. The decomposition uses a logarithmic (log-ratio) approach, which allows the two components to sum exactly to total variance in log terms with no residual. When price and efficiency pull in opposite directions for a given episode, per-episode attribution percentages can exceed 100% or go negative for one side. This reflects the interaction of the two forces and is mathematically correct.

What this study does not claim

The study does not assert that low-cost care groups are, in every case, delivering better clinical care. What is well established is that higher cost does not produce better outcomes.

Dollar attribution figures are specific to the four-metro Texas data set. The proportional finding (efficiency dominates, especially for conditions) is consistent with national research and expected to generalize, but specific dollar magnitudes are Texas-specific.

Physician and patient preference percentages derive from Medicare-based studies. Commercial populations may differ in exact percentages, but the directional findings are broadly consistent across payer types.

Oncology episodes are excluded due to risk adjustment complexity. They are expected in the next iteration.

About the research

The analysis was conducted by Oxbridge Health, which develops episode-of-care payment programs and has launched Episode Benefit Plans in initial markets, supported by reinsurers. These programs are administered by independent TPAs and are available to self-funded employers. The Oxbridge team has previously built and deployed bundled payment models for Medicare and commercial payers.

The peer-reviewed literature cited draws on studies published in Science, JAMA, JAMA Network Open, JAMA Internal Medicine, the New England Journal of Medicine, the Quarterly Journal of Economics, the American Economic Journal: Economic Policy, Health Affairs, the Annals of Internal Medicine, the Journal of Health Economics, Medical Care, PLoS ONE, and other journals, as well as NBER working papers and the Dartmouth Atlas of Health Care.