

# Medtech's custom-build default has outlived its rationale

## Why medtech and diagnostics need a new commercial technology investment thesis

October, 2026 | *By Dylan Herman-Dunphy, SVP, Slipstream Life Sciences, 15+ years helping life sciences companies modernize commercial engagement*

---

By the time a morning surgical case ends, the implant has been placed, its identifying information has been recorded, and the field representative is preparing for the next account. Yet the transaction may remain open for days. Before the manufacturer can issue an invoice, the products documented as used during surgery may need to be reconciled against field inventory records and the hospital's purchase order. Each step depends on information captured by different people, often in separate systems. The resulting delays are not necessarily the consequence of any single failure, but of processes that evolved independently and were never fully integrated.

Many medtech and diagnostics organizations still operate on a long-standing belief: **our operating model is too specialized for standard commercial software, so we have to build.** That belief was reasonable when it formed, and the custom build it produced has since become the commercial system of record by accumulation rather than by decision. The condition that justified it no longer holds.

October 2026

### EXECUTIVE SUMMARY

---

The original reasoning was sound. Medtech and diagnostics revenue depends on a completed operating event rather than a prescribing decision, and for years no horizontal platform supported that, so IT teams extended the CRM, ERP, order and contracting systems they already had. Those extensions became mission-critical, and the technical debt they carry now holds back the operations they were built to support.

What has changed is the technology supply side. Enterprise platforms now provide materially more of the capability these operating models require out-of-the-box.<sup>[6][7][8]</sup> Regulators have moved in the same direction: FDA's Computer Software Assurance guidance now directs a risk-based approach that lets manufacturers rely on supplier evidence and targeted regression testing rather than uniform revalidation.<sup>[11]</sup> Standard functionality can draw on that evidence and custom code cannot, which tilts the assurance burden toward standard capability.

The question is therefore no longer whether each workflow can be custom-built, but which common controls should now be standardized, which legitimate variation should be configured, and which differentiated capability still justifies proprietary development.

## Point of view

Medtech and diagnostics should treat commercial technology as an operating control system, not as a sales productivity tool. Our analysis of public filings shows device companies spend 40% of revenue on SG&A against 23% in pharma and carry more than twice as much inventory relative to revenue (Section 3, with methodology in Appendix A). Those two lines are where better operating control pays back.

## Four decisions to consider before the next platform commitment

1. Price the default: Quantify what customization costs to run and change, separately from what it delivers. Most organizations can say what their customizations deliver but not what they cost, which is why the decision to keep building is rarely challenged.
2. Map the operating model before evaluating platforms: Test native capability against the revenue events your business actually runs, not against a consolidated requirements list.
3. Settle the integration pattern before selecting a vendor: Decide whether the field event will be published once, as a governed contract that downstream systems subscribe to, because that choice constrains platform selection rather than following from it.
4. Allocate customization deliberately: Standardize common industry controls, configure legitimate variation, and build only where a capability creates measurable advantage or native support is still immature, with the burden of proof on whoever proposes to build.

## In this paper we will cover

- |                                          |                                         |
|------------------------------------------|-----------------------------------------|
| 1 The cost of customization              | 7 Standardize, configure or build       |
| 2 Operating model                        | 8 The investment case                   |
| 3 What public filings show               | 9 The final decision and where to start |
| 4 From financial signal to control       | A Appendix: Methodology and statistics  |
| 5 Diagnostics: two businesses, one label | B Appendix: Sources and citations       |
| 6 What has changed in technology supply  |                                         |

## 1. THE COST OF CUSTOMIZATION

### What the custom build actually costs, and how to measure it

We start with cost, the first of the four decisions, because it is the part of the custom-build default that organizations see least clearly, and every later decision depends on it. Customization creates value when it encodes a differentiated capability and becomes overhead when it recreates a common control or makes the platform harder to change. That distinction is rarely audited, because the cost is spread across budgets that are seldom consolidated. Four costs matter most.

## Change cost

Every custom workflow expands the regression surface. In a validated environment each release must show that the modified path still behaves as documented, so routine upgrades become projects, organizations defer them, and the risk compounds as the gap widens. Heavily modified data models also slow adoption of new native capability, because a feature the organization has already built around is only usable once someone can unwind the workaround. That is how a decision made in 2012 still governs what is possible in 2026.

This cost is also easy to measure. Releases behind current, read against the vendor's support policy, take a day to count, are auditable, and past a certain point become a security and compliance exposure rather than a roadmap problem. It is the most board-legible technology risk metric in a regulated environment, and almost no organization reports it.

## Data cost

Each point solution adds interfaces and reconciliation logic, and often creates a second version of the same transaction. A usage event may exist in the field application, the inventory system and the billing extract with three timestamps and no authoritative version. The reconciliation that follows never appears in a technology budget, because commercial operations and finance absorb it.

## People cost

A heavily customized environment concentrates knowledge. The people who can safely change a fifteen-year-old integration are few, and the only people who can assess the cost of changing it are often the people whose work created it. Practitioners observe that critical knowledge in medtech commercial organizations tends to live in people rather than systems, and leaves when they do.<sup>[12]</sup> Engineering capacity meanwhile goes to maintaining common functions that are now available natively, rather than to installed-base analytics or faster usage-to-cash.

## Foreclosed capability

The last cost looks forward rather than back. Governed data, consistent permissions and traceable lineage are preconditions for automation, advanced analytics and AI in commercial and operational workflows, and a fragmented custom environment lacks all three. The cost of the custom-build default is therefore measured in options as well as money, and where these tools touch production or the quality management system, FDA now treats that governance as a regulatory expectation (Section 6).<sup>[11]</sup>

The original decision to custom build on enterprise platforms was right at the time, because purpose-built options for medtech and diagnostics did not exist. **That market has changed, and organizations can no longer afford to run on a default set for conditions that no longer hold.**

## Measure the custom build before debating it

We think the cost of the custom build has to be made visible before anyone debates platforms. Pricing the default is where most organizations stall, because the cost sits in several budgets and nothing aggregates it. The baseline below can be produced in about two weeks from systems the organization already runs. It is deliberately countable: the goal is not precision but replacing an assumption with a number that finance and IT can both see.

| Measure                      | What to count                                                                                                                               | Where it comes from                       |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Configuration mass</b>    | Custom objects and custom fields on core commercial objects, and the share not read or written in the last twelve months                    | Platform metadata and field usage reports |
| <b>Integration surface</b>   | Interfaces by pattern (batch, point-to-point, event) and by owning team, plus the count of transactions represented in more than one system | Integration inventory and middleware logs |
| <b>Version lag</b>           | Releases behind current and standing against the vendor support policy                                                                      | Vendor release notes and support matrix   |
| <b>Bus factor</b>            | Named individuals who can safely change each critical integration, and how many of them are contractors                                     | Team mapping and vendor contracts         |
| <b>Run and change split</b>  | Internal FTE plus vendor spend on maintaining custom functionality, separated from spend on new capability                                  | Finance and resource plans                |
| <b>Shadow reconciliation</b> | Hours per month spent in commercial operations and finance reconciling the same transaction across systems                                  | Process observation and time sampling     |

*Figure 1. Custom-build baseline. Custom code volume, release friction and incident attribution can extend it.*

**Shadow reconciliation** deserves particular attention. It is the only measure in the baseline that sits in a finance budget rather than a technology one: labor cost in commercial operations and finance that a control improvement can remove directly. That makes it the number most likely to turn a technology argument into a finance argument.

Cost alone does not settle the question, though. Before deciding what to standardize, an organization has to understand why it built in the first place, and that comes down to how medtech and diagnostics make money. That is the second decision, and the subject of the next section.

## 2. OPERATING MODEL

### Pharma often sells a decision. Medtech and diagnostics monetize an operating event.

Medtech and diagnostics organizations did not customize by accident. Their commercial model differs from pharma's in ways any replacement has to support. In retail pharma the commercial act typically ends in a prescribing decision, with fulfillment through a pharmacy or specialty channel. Across much of medtech and diagnostics, revenue instead depends on a completed operating event: the manufacturer may need to stage inventory, support the procedure or test, and reconcile what happened before it can bill, so the cost of serving the customer begins when an opportunity closes rather than ending there.

### Five operating archetypes, five technology spines

Each archetype has a different revenue event and therefore needs a different technology spine underneath it. Most organizations of scale run several at once, and a consolidated requirements list averages them into a single specification that matches none. That is why the second decision calls for testing native capability against each archetype's revenue event rather than a merged list.

| Archetype                                          | Revenue event                                                                  | Operating burden                                                                                  | Technology spine                                                                      |
|----------------------------------------------------|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Procedure-intensive and implantable devices</b> | A scheduled case with documented usage                                         | Case coverage, trunk and consignment inventory, lot and serial traceability, usage reconciliation | Case management, field inventory, account hierarchy, order-to-cash integration        |
| <b>Capital equipment and service</b>               | Placement, then uptime, service and eventual replacement                       | Installed-base accuracy, service entitlement, parts, contracts and renewals                       | Installed-base master, field service, contract lifecycle, asset telemetry             |
| <b>IVD instruments and consumables</b>             | Placement followed by reagent pull-through                                     | Menu adoption, utilization, reagent forecasting, service and multiyear contracts                  | Placement economics, consumption, forecasting, service and contract data              |
| <b>Laboratory services</b>                         | An order, a completed test, a covered claim and a collected payment            | Coverage, coding, prior authorization, medical necessity, denial and appeal                       | Referral and account layer plus reimbursement and revenue-cycle data                  |
| <b>Connected devices and wearables</b>             | A device sale, subscription, monitoring relationship or outcome-linked service | Identity, consent, device data, engagement, service and recurring billing                         | Customer and device identity, usage metering, service, subscription and data controls |

## One event, several decision-making authorities

The commercial organization also sells to more than one authority. A clinician may prefer the product while an institutional buyer controls access and a payer controls reimbursement, each on its own cycle and evidence standard, a separation that US value-analysis structures and European tendering formalize rather than create.<sup>[3][4]</sup> A call plan and an opportunity record cannot represent that transaction. The commercial system has to connect demand to the institution that authorizes purchase, to the product or asset that creates revenue, and to the rules that determine payment and follow-on obligation.

**A single field event can change inventory, billing, service and quality records at the same moment, and everything else in this paper follows from that fact.**

## 3. FINANCIAL ANALYSIS

### Public filings put a price on the medtech operating model

If the operating model really is different, it should show up in the numbers. We wanted to test whether public financial data supports that view, knowing it can offer a hint rather than proof: filings do not reveal how technology is built, but they do show where these businesses spend money and tie up capital. We analyzed 382 company-year records across 70 US filers from fiscal 2019 through 2024, above a \$500 million annual revenue floor. Each company is reduced to one median observation so repeated years do not inflate the sample, and mixed-portfolio filers are excluded. The medtech column is device companies only; diagnostics are analyzed separately in Section 5.<sup>[1][2]</sup> The headline results are summarized on the next page.

| Company-level median, FY2019 to FY2024 | MedTech (devices) | Pharma       | Assessment                       |
|----------------------------------------|-------------------|--------------|----------------------------------|
| SG&A as % of revenue                   | 40.0% (n=29)      | 23.0% (n=18) | About 1.7x. Broad and consistent |
| Inventory as % of revenue              | 19.3% (n=29)      | 7.1% (n=18)  | About 2.7x. Real but more modest |

Figure 2. Company-level medians. Both differences are statistically significant; the expense-structure difference is the stronger of the two. Full statistics, sample reconciliation and robustness checks are in Appendix A.

Two objections deserve a direct answer, and Appendix A documents both. First, the sectors split cost of goods and operating expense differently; combining them gives a total cost to serve a dollar of revenue of 78.7% in medtech against 46.1% in pharma, a larger gap than SG&A alone. Second, period-end inventory days do not separate the sectors because pharma's high gross margins inflate the day count, which is why the paper uses inventory as a share of revenue, the quantity a working-capital release actually acts on. Neither SG&A nor inventory intensity is a proxy for technology opportunity, but together they show that commercial execution and product availability are financially consequential enough in these businesses to justify company-specific control and a real investment case.

#### 4. FROM SIGNAL TO CONTROL

##### Translating a financial signal into a technology conclusion

An expense ratio cannot justify a platform decision on its own, and finance is right to reject business cases that treat it that way. **A high SG&A ratio becomes a technology problem only when a specific operating burden can be traced to a specific missing control.** The table below makes that connection for the four mechanisms that matter most, with the measure each function would use to judge the outcome.

| Signal and mechanism                                                                                                                                             | Control layer                                                                                                   | Finance measure                                                                                    | Technology measure                                                                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| High inventory relative to revenue. Product sits across trunks, consignment sites, distributors and customer locations, with demand tied to cases or placements. | A common data model for product, location, lot, serial, expiry, transfer, projected demand and documented use.  | Addressable inventory, expiry and loss, transfer time, case fill rate, usage-to-invoice latency.   | Systems of record consolidated for lot and serial data. Count of inventory reconciliations retired.        |
| High SG&A relative to revenue. Execution requires coordination across sales, operations, service, finance, quality and channel partners.                         | Shared account, asset, contract and workflow records, with fewer custom reconciliations and point integrations. | Manual touches, cycle time, support cost, time to change a territory or launch.                    | Custom objects retired, interfaces removed, cost per release, regression hours.                            |
| Divergent diagnostics models. IVD depends on installed base and consumption. Laboratory services depend on coverage and payment.                                 | A shared customer layer with genuinely different transaction spines above it.                                   | Placement utilization and pull-through, or authorization turnaround, denial rate and days to cash. | One customer model maintained rather than two. Duplicate account masters eliminated.                       |
| Regulated field execution. A single field event can create inventory, billing, service and complaint obligations simultaneously.                                 | Event-driven integration, role-based control, audit history and quality escalation.                             | Completeness of usage capture, traceability, complaint routing, reconciliation exceptions.         | Event contracts published, point-to-point interfaces retired, assurance evidence reusable across releases. |

Figure 3. Signal to control. The finance and technology columns are the two stakeholders' languages, and each row is the contract between them.

## Two chains worked through

### Field inventory

The signal is inventory at 19.3% of revenue. The mechanism is that product sits across locations the manufacturer does not control, and consumption is proven by an event that happens in an operating room rather than in a system. The missing control is rarely a warehouse system; more often it is the absence of an authoritative link between a documented case and an inventory movement. The right measures are addressable inventory by location and status, expiry and loss against baseline, and case fill rate held constant while stock falls. If fill rate degrades, the program has destroyed value whatever the balance sheet shows.

### Usage-to-invoice latency

The signal is SG&A intensity, but the mechanism is more specific than headcount. Between a completed procedure and a submitted invoice sits a reconciliation process that, in most organizations, is manual, distributed and undocumented. The missing control is a single usage event that operations and finance both accept as authoritative. The measures are elapsed days from event to invoice, the count of unbilled completed events at period end, and hours spent on reconciliation exceptions. Each can be baselined in weeks, which makes this the most commonly winnable case and the natural first increment discussed in Section 8.

### Integration pattern precedes platform selection

Because one completed case can update inventory, billing, service entitlement and a potential complaint record, point-to-point integration between the commercial system and each downstream system will not hold. Every new obligation multiplies interfaces, and every interface is another place where the same transaction acquires a second version, a consequence that is usually discovered too late.

This is the third decision: whether to publish the field event once, as a governed contract that downstream systems subscribe to, and it should be settled before vendor evaluation begins. An organization that skips it will reproduce its current integration topology on whatever platform it buys.

Purpose-built technology does not remove the underlying complexity of these chains, but it reduces the custom machinery required to represent and control them. That, rather than a generic productivity claim, is the economic basis for reconsidering the build default.

## 5. DIAGNOSTICS

---

### Diagnostics is two businesses under one label

We treat diagnostics separately because it shares medtech's core challenge, revenue tied to an operating event, but behaves very differently underneath. It is usually bought as one segment and staffed as one commercial organization, yet it contains two businesses whose inventory intensity differs by a factor of three and a half. **An organization that cannot distinguish them inside its own category will write one requirement, buy one platform and serve neither well.**

The difference is invisible on the metric most organizations check first, because IVD manufacturers and laboratory-service providers have effectively identical SG&A intensity. The balance sheet tells a different story.

| Company-level median, FY2019 to FY2024 | IVD manufacturers | Laboratory services | Assessment                      |
|----------------------------------------|-------------------|---------------------|---------------------------------|
| Inventory as % of revenue              | 13.0% (n=7)       | 3.7% (n=6)          | 3.5x. Statistically significant |
| Period-end inventory-equivalent days   | 118 (n=6)         | 36 (n=5)            | 3.3x. Statistically significant |
| SG&A as % of revenue                   | 28.7% (n=7)       | 26.0% (n=7)         | No meaningful difference        |

Figure 4. Small samples with very large effects, treated as directional. The public IVD universe is limited because the largest participants are foreign private issuers or segments of diversified filers. Full statistics are in Appendix A.

The difference matters for the build decision. IVD manufacturers earn revenue through instrument placement and reagent pull-through, so their commercial system needs installed-base and consumption data; without it, problems surface as missed forecasts and surprise renewals. Laboratory services earn revenue only when a test moves from order to paid claim under coverage rules such as CMS MoIDX,<sup>[5]</sup> so their commercial system has to reach into the revenue cycle. **The customer layer can be standardized across both businesses, while each transaction spine needs its own configuration.** Organizations that miss this tend to write one generic diagnostics requirement and then custom-build their way around it when the second business tries to go live.

## 6. WHAT HAS CHANGED IN THE TECHNOLOGY SUPPLY MARKET

### The platform gap has narrowed, and regulation has moved in the same direction

So far we have argued that the custom build is expensive and that the operating model behind it is real. This section explains why the default no longer holds: independent sources describe the same accumulation of customization, the platform market has closed much of the gap, and the regulator now favors standard software.

### Further evidence of the legacy burden

The financial analysis in Section 3 offers a hint rather than direct evidence, so it matters that other sources point the same way. Recent practitioner analysis describes medtech commercial data fragmented by years of platform proliferation, CRM customization and organizational silos, with procedure data in a case management system, pricing and contracts in a revenue management platform, inventory in ERP and interaction history in a CRM configured differently by each business unit.<sup>[12]</sup> McKinsey's 2022 survey of more than 1,900 medtech leaders, the largest sample cited here, found fewer than half considered their technology architecture reliable and only 17% reported a modular one.<sup>[10]</sup> A 2024 vendor-sponsored benchmark of more than 130 commercial leaders found nearly 70% rely on manual processes and homegrown applications to manage commercial content and claims, and only 15% link claims directly to substantiation.<sup>[9]</sup> The burden these sources describe is an accumulation of homegrown applications and customization, not an absence of software.

### What the platform market now provides

Current enterprise offerings include capability that was previously absent or routinely custom-built, and official product documentation now describes native support for procedure and inventory workflows, richer account structures and commercial engagement planning. The offerings are not interchangeable, none eliminates integration with core operational systems, and none is recommended here. The point is only that the gap that once justified building is smaller than it was.<sup>[6][7][8]</sup>

What the regulator now expects

The assurance economics have moved in the same direction, and this is the development most often missed. FDA’s final guidance on Computer Software Assurance, issued September 2025 and updated February 2026 to align with the Quality Management System Regulation, directs a risk-based approach keyed to intended use and process risk.<sup>[11]</sup> It permits manufacturers to rely on supplier certifications and development documentation as assurance evidence, prefers system logs and audit trails over paper, and calls for targeted regression testing rather than full revalidation when software-as-a-service functionality updates. Its scope includes AI and machine-learning tools used in production or the quality management system.

The resulting asymmetry is the point. Supplier evidence exists for standard vendor functionality and not for code the organization wrote itself, so two systems delivering the same control no longer carry the same assurance burden. The guidance does not govern commercial CRM as such, but it applies wherever the commercial system reaches into the quality record, which in these operating models is a recurring boundary rather than an edge case.

What to test before selection

The gap has narrowed rather than closed, so how a platform is tested matters as much as which one is chosen. Demonstrations show where native capability is strongest; test where it most often stops short: hierarchy governance, stewardship and indirect channels in the account model; offline operation, ERP integration and recall and complaint linkage in field inventory; system-of-record boundaries, telemetry and parts integration for installed base and service; calculation depth, auditability and regional procurement in contracts; regional compliance and content integration in customer engagement; and lineage, permissions, assurance evidence and human controls for AI and automation. Run each test against a real revenue event from the archetypes in Section 2, as the second decision calls for, not against a feature checklist.

7. ARCHITECTURE DECISION

Standardize, configure, or build

A narrower platform gap and a regulator that favors standard software together move the burden of proof onto whoever proposes to build, which is the fourth decision. The conclusion is not to stop building, but to allocate custom development deliberately rather than by default. Reproducing every legacy workflow on new technology preserves the constraints that produced it, while replacing every differentiated workflow with a standard process destroys what the company competes on. Both failures come from making the decision at program level, when it belongs at capability level.

| Requirement type                              | Preferred treatment | Decision test                                                                                         |
|-----------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------|
| Common industry capability supported natively | Standardize         | Does the native process satisfy regulatory and operating requirements without material loss of value? |
| Legitimate company or regional variation      | Configure           | Can metadata, rules, roles or workflow configuration represent the variation without custom code?     |
| Differentiated commercial capability          | Build selectively   | Does the capability create measurable advantage that a standard process would remove?                 |

| Requirement type                                           | Preferred treatment      | Decision test                                                                          |
|------------------------------------------------------------|--------------------------|----------------------------------------------------------------------------------------|
| <b>Legacy process created by system limitations</b>        | Redesign                 | Would the company choose this workflow if it were designing the operating model today? |
| <b>Important requirement with immature product support</b> | Extend with an exit plan | Can the extension be isolated, governed and retired when native capability matures?    |

## Applying the test honestly

Two failure modes turn this table into decoration. The first is classifying a legacy workaround as differentiated capability, which is how organizations rebuild their constraints while believing they are protecting their advantage; the fourth row exists to catch it. The second is assuming native means complete, when purpose-built products can be incomplete, expensive or overly prescriptive and migration carries real execution risk. An extension with a documented exit plan is a legitimate answer, and one without is how the next custom build begins.

## 8. THE INVESTMENT CASE

### Build the business case on a first measurable outcome

Knowing what to standardize, configure and build still leaves the question every CFO will ask: is the change worth it? We think the answer should be proven in increments rather than asserted up front. Programs that move away from the custom-build default most often fail not by selecting the wrong platform but by committing to a multi-year replacement whose first measurable outcome sits eighteen months out, by which point the sponsor has changed and the business case can no longer be tested.

The fastest place to prove it is usually the usage-to-invoice chain in Section 4, which can show a measurable result within two quarters, before a full platform decision is made. Whatever the starting point, the full case should pass the test below.

#### The investment test

A purpose-built platform is justified when expected operating value plus avoided custom-build cost exceeds the full cost and risk of change, within the company's own payback threshold.

**Count the whole cost of change:** license, implementation, assurance and re-qualification, data migration, parallel run, change management, and the deferred upgrade backlog inherited on day one. Migration carries its own assurance burden, and a business case that omits it will not survive finance review.

**Keep the benefits honest:** a working-capital release is a one-time benefit whose recurring value is the avoided carrying cost, and faster billing is not incremental revenue unless the operating change also increases demand or improves realization.

## 9. THE FINAL DECISION

### The market that justified custom building has changed, and the build default should change with it

Every workflow could be rebuilt on new technology; that is how today's custom footprint came to exist. **The real question before the next platform commitment is which common controls should now be standardized, which variation deserves configuration, and which capabilities still justify proprietary development.**

Organizations that answer it before a renewal or transformation can reduce the cost of customization without giving up differentiation. Those that do not will reproduce their legacy customization on newer technology, at higher cost, and call it modernization.

**Where to start:** Slipstream helps medtech and diagnostics organizations test the custom-build default against their own data. In four to six weeks, we analyze your current commercial technology footprint and deliver a roadmap to simplify it: where to retire customization, how to reduce technical debt and the cost of maintaining it, and how to build a scalable platform that drives business value. To connect with Slipstream reach out at: <https://www.slipstreamls.com/contact/>

---

## ABOUT SLIPSTREAM LIFE SCIENCES

---

Slipstream Life Sciences is a technology consultancy focused exclusively on life sciences, supporting 240+ clients with 700+ resources. We help medtech and diagnostics organizations translate distinctive operating models into scalable commercial platforms and data foundations.

Our work covers the full platform lifecycle, from strategy and architecture through implementation and ongoing optimization, with a focus on what makes these businesses distinct: field execution, installed-base economics and reimbursement-driven workflows. The objective is to reduce unnecessary customization without removing capability that creates value.

*Disclosure: Slipstream Life Sciences holds partnerships with commercial platform vendors, including vendors whose product documentation is cited in Appendix B. No provider is named or evaluated in the body of this paper, and no provider sponsored or reviewed it.*

## ABOUT THE AUTHOR

---

Dylan Herman-Dunphy, SVP, Slipstream Life Sciences, has 15+ years helping life sciences companies modernize commercial engagement where he has built three practices from the ground up and leads a global team of 150+ architects, engineers and consultants.

During his time in healthcare and life sciences technology consulting, he has designed and delivered enterprise platforms across commercial, clinical, medical, marketing, patient services and manufacturing for medtech, diagnostics, pharma and research organizations.

His recent work sits where AI meets compliance. His team has built agentic solutions in Agentforce. He brings hands-on experience with Google Vertex AI, and leads the team that has built ElevenLabs voice agents and established some of the first life sciences use cases for them.

He also leads Slipstream's OneTrust privacy practice, helping clients govern consent and data use under HIPAA, GDPR and emerging state privacy laws. He partners with senior executives to turn complex operating models into platforms that scale.

# APPENDIX A

## Financial analysis methodology and statistical results

### Dataset and unit of analysis

The dataset contains 382 company-year records across 70 public-company filers for fiscal years 2019 through 2024. Each included company-year meets a \$500 million revenue floor. Headline tests are conducted at company level: each company is reduced to its median across available years, then segment distributions are compared. This prevents repeated years from the same company inflating the effective sample size. Comparisons use the two-sided Mann-Whitney U test and Cliff's delta. No multiple-comparison correction is applied; the headline results survive any standard correction, and the diagnostics results are reported as directional.

### Sample reconciliation

| Segment             | Filers | Company-years | Exclusions applied                                              | Headline n |
|---------------------|--------|---------------|-----------------------------------------------------------------|------------|
| MedTech (devices)   | 31     | 167           | Mixed-portfolio: BDX, BAX                                       | 29         |
| Pharma              | 19     | 108           | Animal health: ZTS                                              | 18         |
| IVD manufacturers   | 13     | 74            | Mixed-portfolio: TMO, DHR, A.<br>Tools-adjacent: BRKR, WAT, MTD | 7          |
| Laboratory services | 7      | 33            | None                                                            | 7          |
| Total               | 70     | 382           | Section 3 headline uses 29 + 18                                 | 47         |

*Figure A1. The Section 3 headline compares 47 of the 70 filers. Metric-specific n differs where a component is unavailable; those counts are reported in every table.*

### Headline filters and definitions

Mixed-portfolio filers are excluded from segment medians where consolidated reporting impairs comparability.

Tools-adjacent filers are excluded from the IVD headline and retained as a sensitivity set. Their medians for reference: inventory 14.3% of revenue, 124.5 inventory-equivalent days, SG&A to R&D 3.8.

Zoetis is excluded from the pharma headline because animal health has a different commercial and inventory model. Including it narrows the medtech-pharma differences slightly but creates a conceptual peer-set weakness.

**Ratios are computed per company and then medianed:** consequently the SG&A to R&D figures will not reconcile against the ratio of the SG&A and R&D medians in the same table. Ratio of medians is 5.39 for medtech devices and 0.97 for pharma, against the reported medians of ratios of 5.58 and 0.98. Point differences are computed on unrounded values.

Period-end inventory-equivalent days equals period-end inventory divided by annual cost of goods sold, multiplied by 365. It is not average-inventory DIO.

Cost of goods sold has the weakest coverage of any field, at 45 nulls across 382 records, which is why measures that divide by it carry lower n. Metric-specific sample sizes are reported throughout, and missing facts are not converted to zero.

**Incyte figures were sourced manually:** automated XBRL extraction returned revenue but no other fields for this qualifying pharma filer, so its figures were taken directly from the company's SEC filings. Inventory is period-end total inventory, current plus non-current, from the consolidated and condensed consolidated balance sheets: \$16.5M at FY2019, \$36.0M at FY2020, \$56.9M at FY2021 and \$121.0M at FY2022. Income statement figures were taken from reported annual results. The robustness checks show the effect of excluding it.

## Primary statistical results

| Metric                      | MedTech (devices) | Pharma       | Mann-Whitney p | Cliff's delta |
|-----------------------------|-------------------|--------------|----------------|---------------|
| SG&A as % revenue           | 40.0% (n=29)      | 23.0% (n=18) | 0.0001         | +0.69         |
| R&D as % revenue            | 7.4% (n=29)       | 23.7% (n=18) | <0.001         | -0.76         |
| SG&A to R&D                 | 5.58 (n=29)       | 0.98 (n=18)  | <0.001         | +0.88         |
| Inventory as % revenue      | 19.3% (n=29)      | 7.1% (n=18)  | 0.0015         | +0.56         |
| Inventory-equivalent days   | 172 (n=28)        | 174 (n=17)   | 0.842          | +0.04         |
| Gross margin                | 63.9% (n=28)      | 78.3% (n=17) | 0.0002         | -0.68         |
| COGS plus SG&A as % revenue | 78.7% (n=28)      | 46.1% (n=17) | <0.001         | +0.84         |

*Figure A2. Cliff's delta ranges from -1 to +1. Positive values mean a randomly selected medtech company is more likely to record a higher value than a randomly selected pharma company. Note the material difference in effect size between the expense-structure findings and the inventory finding; both are significant, they are not equally strong. The inventory-equivalent-days result is effectively zero and is reported rather than omitted.*

## Decomposition of the SG&A to R&D ratio

Because the ratio has a denominator that differs sharply across sectors, its interpretation is narrower than its effect size suggests. Medtech's SG&A intensity is 1.74 times pharma's. Pharma's R&D intensity is 3.19 times medtech's. Approximately 68% of the logarithmic gap in the ratio is attributable to the R&D denominator. The ratio is therefore evidence that spend concentrates differently, not that commercial cost alone separates the sectors. SG&A as a percentage of revenue is the load-bearing measure for the latter claim.

## Allocation-neutral robustness

Combining cost of goods sold with SG&A removes any dependence on where the two sectors draw that line. The combined measure produces a larger effect than SG&A alone (delta 0.84 against 0.69) and both components are individually higher in medtech: roughly 36% COGS plus 40% SG&A for medtech devices, against roughly 22% plus 23% for pharma. This measure is not a commercial-cost measure and is used only to establish that the SG&A finding is not a classification artifact.

## Choice of inventory measure

Period-end inventory-equivalent days are effectively identical across medtech and pharma (172 against 174). Days divides by cost of goods sold, so pharma's high gross margins produce high day counts even where inventory is light relative to revenue, which makes it a poor cross-sector comparison. Inventory as a percentage of revenue measures capital committed per dollar of revenue, the quantity a working-capital release actually acts on. Within diagnostics both measures agree (Section 5).

## Diagnostics segmentation

| Metric                    | IVD manufacturers | Laboratory services | p     | Cliff's delta (95% CI) |
|---------------------------|-------------------|---------------------|-------|------------------------|
| Inventory as % revenue    | 13.0% (n=7)       | 3.7% (n=6)          | 0.002 | +0.95 (0.71 to 1.00)   |
| Inventory-equivalent days | 118 (n=6)         | 36 (n=5)            | 0.009 | +0.93 (0.60 to 1.00)   |
| SG&A as % revenue         | 28.7% (n=7)       | 26.0% (n=7)         | 0.620 | -0.18 (-0.76 to 0.47)  |

*Figure A3. Confidence intervals are bias-corrected percentile bootstrap, 10,000 resamples. Despite n of 7 and 6, both inventory intervals remain well clear of zero, while the SG&A interval correctly spans it. The public IVD universe is structurally constrained because the largest participants are foreign private issuers or segments within diversified filers, and cannot be isolated from 10-K XBRL. Treated as directional evidence supporting an operational distinction, not as a precise industry estimate.*

## Pooled medtech and diagnostics

Section 3 reports device companies separately. Pooling all three non-pharma segments against pharma gives SG&A of 35.3% of revenue (n=43) against 23.0%, p=0.002, delta +0.50, and inventory of 14.8% of revenue (n=42) against 7.1%, p=0.020, delta +0.38. Both differences remain significant, though the inventory effect weakens by about a third. Devices are reported separately because pooling blends the IVD and laboratory-services segments, whose inventory intensity differs by roughly three and a half times (Section 5).

## Robustness checks

We include these checks because the financial analysis is a hint rather than proof, and a reader should be able to see that the hint is not an artifact of how the sample was built. Each changes one sample construction choice; figures are medtech against pharma. A balanced six-year panel (companies with six non-null years for the metric tested) gives SG&A to R&D of 5.85 against 1.12, inventory of 20.3% against 12.5% of revenue, and combined cost to serve of 76.6% against 46.8%. Restricting to FY2022 to FY2024 gives 5.62 against 0.97, 20.8% against 7.5%, and 78.1% against 46.7%, so the pattern is not solely a pandemic-period effect. Requiring at least three years per company gives 5.51 against 0.98, 19.2% against 7.1%, and 76.6% against 46.1%. A restricted peer set leaves SG&A to R&D and inventory relative to revenue in the same direction and of similar magnitude. Excluding Incyte, the one manually sourced filer, moves pharma SG&A to R&D from 0.98 to 1.06 and inventory from 7.1% to 7.6%, with slightly smaller effect sizes on both.

## Limitations

SG&A is broader than commercial expense. Inventory relative to revenue describes balance-sheet intensity, not avoidable stock. The combined cost measure includes manufacturing and is not a commercial-cost proxy. The sample is limited to US 10-K filers above a \$500 million revenue floor, which excludes the largest medtech and diagnostics participants where they are foreign private issuers or segments within diversified filers, and may therefore skew toward mid-cap pure-plays. The analysis does not observe technology architecture, process quality or implementation outcomes, and therefore cannot estimate software return from public filings. Its valid use is to identify the economic control areas that merit company-specific baselining.

# APPENDIX B

## Sources and citation notes

---

Each entry gives the source, its location and the reason it is used. Third-party and vendor-sponsored material is identified as such and was not independently verified. All financial figures derive from sources [1] and [2], except the Incyte figures noted in Appendix A. Entries are grouped by category rather than by order of first appearance in the text.

### Primary financial data

[1] U.S. Securities and Exchange Commission. EDGAR Application Programming Interfaces. <https://www.sec.gov/search-filings/edgar-application-programming-interfaces> Use: official documentation for Company Facts and XBRL APIs.

[2] U.S. Securities and Exchange Commission. Financial Statement Data Sets. <https://www.sec.gov/data-research/sec-markets-data/financial-statement-data-sets> Use: official description and limitations of SEC-extracted financial statement data.

### Commercial model and purchasing authority

[3] IQVIA. Considerations for Successful MedTech Product Adoption. <https://www.iqvia.com/locations/united-states/library/white-papers/considerations-for-successful-medtech-product-adoption> Use: structural context for institutional purchasing, evidence and adoption. Third-party commercial publisher.

[4] European Union. Directive 2014/24/EU on public procurement. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32014L0024> Use: primary source for price-quality procurement criteria.

[5] Centers for Medicare & Medicaid Services. MoIDX: Molecular Diagnostic Tests. <https://www.cms.gov/medicare-coverage-database/view/lcd.aspx?lcdid=35025> Use: primary coverage source for test-specific identifiers in applicable claims.

### Platform capability documentation

*Cited to establish that named capability exists in the market. No comparative or evaluative claim is made about any provider, and no provider is named in the body of this paper. Both vendors are identified by name here so that the citations are verifiable and so that neither is treated differently from the other.*

[6] Salesforce. Life Sciences Cloud MedTech Datasheet, FY26. [https://www.salesforce.com/en-us/wp-content/uploads/sites/4/Life-Sciences-Cloud\\_MedTech-Datasheet\\_FY26.pdf](https://www.salesforce.com/en-us/wp-content/uploads/sites/4/Life-Sciences-Cloud_MedTech-Datasheet_FY26.pdf) Use: official documentation of field inventory, surgical case and provider relationship capabilities.

[7] Salesforce. Help documentation, Life Sciences Cloud commercial engagement. [https://help.salesforce.com/s/articleView?id=ind.lsc\\_landing\\_commercial\\_engagement.htm](https://help.salesforce.com/s/articleView?id=ind.lsc_landing_commercial_engagement.htm) Use: official setup and capability documentation for case visits, cycle counting and inventory workflows.

[8] Veeva Systems. MedTech CRM Suite. <https://www.veeva.com/medtech/products/commercial/crm-suite/> Use: official documentation of unified customer data, commercial applications and embedded automation.

## Architecture maturity and market context

[9] Veeva MedTech. 2024 Commercial Benchmark Report. [https://www.veeva.com/medtech/wp-content/uploads/2024/11/MedTech\\_2024-Commercial-Benchmark-Report\\_v1.pdf](https://www.veeva.com/medtech/wp-content/uploads/2024/11/MedTech_2024-Commercial-Benchmark-Report_v1.pdf) Use: vendor-sponsored survey of more than 130 commercial medtech leaders, November 2024. Source for the 70% manual and homegrown applications figure and the 15% claims-to-substantiation figure. Commercially sponsored and identified as such in the text.

[10] McKinsey & Company. How medtechs can meet industry demand for omnichannel engagement. <https://www.mckinsey.com/industries/life-sciences/our-insights/how-medtechs-can-meet-industry-demand-for-omnichannel-engagement> Use: Omnichannel Maturity Survey of more than 1,900 global medtech leaders, conducted July 2022. Source for architecture reliability and modularity figures.

## Regulatory

[11] U.S. Food and Drug Administration. Computer Software Assurance for Production and Quality Management System Software. Final guidance, February 2026, docket FDA-2022-D-0795, superseding the final guidance of 24 September 2025. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/computer-software-assurance-production-and-quality-management-system-software> Use: primary regulatory source for risk-based assurance, reliance on supplier evidence, targeted regression testing for software-as-a-service updates, and scope covering AI and machine-learning tools used in production or the quality management system.

## Practitioner observation

[12] ZS Associates. Medtech CRM strategy: a customer-first operating model, April 2026. <https://www.zs.com/insights/from-fragmented-capabilities-to-customer-first-medtech-crm> Use: third-party practitioner observation that critical commercial knowledge resides with individuals rather than systems, and that commercial data is fragmented across case management, revenue management, ERP and CRM systems. Cited as independent corroboration in Sections 1 and 6.