

DECISION SCIENCES INSTITUTE

Diagnostic AI and dynamic absorption in sequential care lines

ABSTRACT

Diagnostic AI is often evaluated as a local improvement in accuracy. In sequential care lines, however, new information changes output only if it is produced where downstream action is decided. We study FFR_{CT} adoption in coronary care using a 2016–2024 Medicare hospital-year panel. Adoption expands coronary CT volume by about 43%, but only 6%–8% of that increase reaches the invasive stage, while treated output rises modestly. The treatment-to-imaging ratio falls and then recovers, forming a J-curve. We interpret this as dynamic absorption of a positioned information-processing capability.

KEYWORDS: Artificial intelligence, Healthcare operations, Information processing theory, Sequential service systems, Difference-in-differences

1. INTRODUCTION

Artificial intelligence (AI) is usually introduced into healthcare as an improvement in information processing. A model reads an image, converts noisy clinical signals into a prediction, and gives clinicians a more scalable diagnostic judgment. In a sequential care line, however, information processing is not valuable simply because it is more accurate or cheaper. It is valuable when it is produced at a position where it can change downstream action. This paper studies that positional problem. We ask what happens when an AI tool makes the front of a care line information-rich, but the decision that determines output is made farther downstream.

The question matters because healthcare AI is expanding fastest in diagnostic settings. Imaging AI has grown rapidly, demand for specialist interpretation continues to increase, and regulatory review of software devices commonly focuses on device performance rather than on the operating consequences of deploying the device inside a care pathway (Sivakumar et al. 2025, Benjamens et al. 2020, Jing et al. 2025, Muehlematter et al. 2023). For hospitals, the relevant operating question is therefore different from the regulatory or clinical validation question. The issue is not only whether the AI produces a useful local assessment. It is whether the new information changes the flow, structure, and output of the line in which the assessment is embedded.

We study this question in coronary care. The focal technology is computed-tomography fractional flow reserve (FFR_{CT}), a diagnostic-gate AI that estimates coronary physiology from an already acquired coronary CT. Clinically, tools of this kind are intended to improve triage and reduce unnecessary invasive testing rather than mechanically increase downstream procedures (Douglas et al. 2015, Fairbairn et al. 2018, Zink et al. 2026). That feature makes the setting especially useful. If downstream treatment does not rise proportionately after adoption, the interpretation is not that the technology failed clinically. The interpretation is operational: an upstream information-processing capability may primarily expand and reorganize diagnostic activity, while treatment output remains governed by a downstream action margin.

The coronary care line provides separately observable stations. Patients enter through coronary CT angiography. Some proceed to invasive coronary angiography (ICA) in the catheterization laboratory, and a subset receive revascularization. We assemble a public Medicare hospital-year panel from 2016 through 2024 and exploit 1,201 staggered hospital adoptions of FFR_{CT}. Because adoption is correlated with rising pre-adoption imaging volume, our level estimates do not rely on naive event-study comparisons. We use estimators designed for staggered adoption under se-

lection on pre-period outcome trajectories and read results as a bracket rather than as a single coefficient.

The empirical pattern has three core steps and one boundary test. First, FFR_{CT} adoption increases front-end coronary CT volume by about 43% in the preferred synthetic difference-in-differences estimate. This is the local release of information-processing capacity at the diagnostic gate. Second, the increase does not transmit proportionately downstream: only about 6%–8% of the front-end increase reaches the invasive stage, and treated output rises only modestly, by about +0.03 to +0.07 log points across selection-robust estimators. Third, the resulting diagnostic tilt is dynamic rather than permanent. The treatment-to-imaging structure index declines at adoption and then recovers over subsequent years, a J-curve consistent with care-line absorption of the new information-processing capability. Finally, a within-line position contrast suggests that the signature depends on where physiological AI is placed: upstream deployment expands the front and tilts line structure, whereas downstream deployment does not generate the same front-end signature.

The paper contributes to information-processing theory by moving from the amount of information-processing capacity to its process position. Classical information-processing theory argues that organizations facing uncertainty must increase information-processing capacity or redesign structures so that information demand and processing capacity fit (Galbraith 1974, Tushman and Nadler 1978). We extend that logic to AI in sequential care lines. AI is not a generic capability floating above the organization. It is a positioned capability embedded at a particular station. Its immediate effect is local: it relaxes uncertainty where it is deployed. System output changes only when that position is close enough to the margin that turns information into action. The J-curve further shows that this translation is dynamic: a line absorbs AI-generated information over time rather than converting it into output immediately.


This framing also clarifies the managerial contribution. A diagnostic-gate AI should not be evaluated only by local accuracy or adoption volume. It should be evaluated by what it does to the line: how much front-end activity it creates, how much of that activity reaches downstream action, and how long the line takes to absorb the new information. In our setting, a tool placed at the gate primarily produces diagnostic abundance and structural rebalancing, not an immediate throughput gain. This is not a claim about patient welfare, which public aggregate data cannot identify. It is a claim about care-line operations: where AI is placed determines whether it appears as information production, action conversion, or delayed absorption.

The rest of the paper proceeds as follows. Section 2 describes the coronary care line and the focal AI positions. Section 3 develops the information-processing mechanism and hypotheses. Section 4 presents the data and empirical strategy. Section 5 reports the main results. Section 6 uses a reduced-form funnel to locate the margins through which the front-end shock transmits. Sections 7–9 present robustness checks, discussion, limitations, and conclusions.

2. THE CORONARY CARE LINE AND THE DIAGNOSTIC-GATE AI

2.1 The Coronary Care Line

We treat coronary care as a sequential information-processing line with three observable stations. The front station is coronary CT angiography (CCTA), which produces noninvasive anatomical information. The middle station is invasive coronary angiography (ICA), which confirms anatomy in the catheterization laboratory. The output station is revascularization, measured primarily through percutaneous coronary intervention and, where relevant, coronary artery bypass grafting. A patient can move downstream only after upstream information makes further work-up worthwhile, so each station transforms uncertainty for the next one.



Notes. Schematic of the three-stage line. Fractional flow reserve is computed at two stations: FFR_{CT} at the gate (upstream, alongside plaque AI) and angiographic FFR in the cath lab (downstream).

Figure 1: The coronary care line and the two AI positions.

The stations differ not only in clinical purpose but also in operating logic. The front station can be expanded relatively cheaply by adding image acquisition, image interpretation, or software-based post-processing. The downstream stations are scheduled, capital-intensive, and staffed by specialized teams. A conventional bottleneck view therefore points to the cath lab and treatment end of the line: if the front becomes larger, managers might expect output to increase only when downstream capacity can absorb the added candidates (Goldratt and Cox 1984). The information-processing view developed below adds a second possibility. The binding margin may not be a capacity level at the treatment stage. It may be the rate at which downstream arrivals are converted into treatment-relevant action.


2.2 The Diagnostic-Gate AI and the Position Contrast

FFR_{CT} is a post-processing software service applied to an already acquired coronary CT. It estimates fractional flow reserve from the CT image and is typically interpreted relative to the clinical threshold of 0.80 (National Institute for Health and Care Excellence 2017). In use, the tool can provide functional information before an invasive procedure is ordered. That makes it a gate AI: it affects the information available when the system decides whether a patient should proceed to invasive angiography.

We date FFR_{CT} adoption from billing under HCPCS 75580 and earlier temporary T-codes 0501T–0504T. A hospital is treated as an adopter from the first year either code family appears. Adoption is staggered from 2018 through 2024 and follows a series of regulatory and payment milestones, including FDA clearance, guideline endorsement, the move to a permanent Category I code, and a subsequent increase in Medicare payment (Figure 2).

The coronary line also offers a position contrast. Physiological assessment is observable at multiple stations. Upstream AI, measured through plaque-characterization software (0623T–0626T), operates at the CT station. Downstream AI, measured through angiographic-FFR software (0523T), operates in the cath lab. Both are information-processing tools, but they serve different uncertainties. Upstream tools affect who is worked up and potentially referred to invasive evaluation. Downstream tools inform action conditional on invasive confirmation. A wire-based, non-AI comparator at the downstream station (93571) helps separate a downstream-position effect from generic cath-lab activity.

We do not re-evaluate FFR_{CT} accuracy. Prior clinical studies show that CT-derived fractional flow reserve improves discrimination of flow-limiting disease relative to anatomical CT assessment and agrees with invasive reference standards (Nørgaard et al. 2014, Koo et al. 2011, Faulder et al. 2024). Trials and registries support CT-first pathways that can defer unnecessary invasive testing



Notes. Bars: new FFR_{CT} adopters per year, capability panel (Table A1). Markers: FDA clearance (2014); NICE guidance (National Institute for Health and Care Excellence 2017); the 2021 ACC/AHA Class 1 recommendation (Gulati et al. 2021); the 2024 Category I CPT code; the 2025 CMS coronary-CT payment increase.

Figure 2: FFR_{CT} adoption and regulatory milestones, 2014–2025.

without worse outcomes (Douglas et al. 2015, The SCOT-HEART Investigators 2018, Fairbairn et al. 2018). Our outcome is different: we observe public aggregate counts of line activity. The data identify how adoption changes flow, structure, and output; they do not identify patient-level appropriateness, safety, or welfare.

3. THEORY AND HYPOTHESES

3.1 AI as a Positioned Information-Processing Capability

Information-processing theory starts from a simple proposition: organizations facing greater uncertainty need greater information-processing capacity or a design that reduces the information they must process (Galbraith 1974, Tushman and Nadler 1978). This logic has been productive in healthcare operations because hospitals are information-intensive service systems. Patients arrive with heterogeneous symptoms, uncertain diagnoses, and uncertain treatment needs; clinical work proceeds by producing, exchanging, and acting on information (Devaraj et al. 2013, Gardner et al. 2014, Dobrzykowski and Tarafdar 2015, Srinivasan and Swink 2018, Thirumalai et al. 2024).

Diagnostic AI appears, at first, to fit directly into this theory. It adds information-processing capacity by making a clinical assessment cheaper, faster, or more scalable. In a sequential care line, however, information-processing capacity has a position. An AI placed at the front of a line processes admission-relevant uncertainty: who should be worked up, who should be referred, and which cases should be ruled out before invasive resources are used. An AI placed near the end of the line processes action-relevant uncertainty: which downstream arrivals should become treatment. The two uncertainties are related, but they are not the same. A tool can be accurate at the upstream task and still be distant from the margin that determines system output.

This distinction extends classical information-processing theory. The issue is not only whether the organization has more information-processing capacity. The issue is where that capacity is embedded and whether the station it serves controls the conversion of information into action. We call this a positioned information-processing capability. Its local effect is to relax the information constraint at its own station. Its system effect depends on whether the line can transmit and absorb that local information downstream.

3.2 From Information Production to Action Conversion

A sequential care line separates information production from action conversion. At the front, additional sensing can enlarge the candidate pool by making work-up cheaper or more informative (Agrawal et al. 2019). Downstream, candidates become output only if they pass through referral and convert into treatment. Thus, an upstream AI must cross two margins before it changes system output. First, additional front-end candidates must reach the invasive stage. Second, invasive arrivals must become treatment.

This separation creates the central prediction of the paper. If the AI is placed at the diagnostic gate, it should first expand front-end information production. The expansion may be valuable for screening and deferral, but it need not increase treatment proportionately because the downstream action margin remains governed by other decisions, routines, and constraints. In the language of prediction economics, cheaper information changes action only when the complementary decision process can use it (Goldfarb and Tucker 2019). In a single clinical decision, prediction and action may be co-located. In a care line, they are separated by referral, scheduling, and treatment evaluation.

3.3 Hypothesis Development

Front-end information production (H1). When sensing is costly, hospitals tend to reserve it for cases likely to matter. When sensing becomes cheaper and more scalable, the front end can be used on a broader population. A diagnostic-gate AI therefore increases the amount of information produced at the station where it is deployed. In coronary care, this means more coronary CT activity after FFR_{CT} adoption. This prediction does not require a claim that every additional scan is appropriate or inappropriate; it is a statement about the operating response to new front-end information-processing capacity.

Hypothesis 1 (H1). *After a hospital adopts FFR_{CT}, upstream coronary CT imaging volume increases.*

Weak immediate information-to-action transmission (H2). Front-end information production is not the same as downstream action. A larger candidate pool should translate into proportional treatment output only if the added information changes the downstream conversion margin. For an upstream AI, that condition is not automatic. The tool helps decide who should be further evaluated, but treatment still depends on downstream clinical confirmation, catheterization decisions, patient selection, and revascularization judgment. Thus, the diagnostic front can expand while treatment output grows more slowly.

The observable implication is a diagnostic tilt in line structure. We measure it with the treatment-to-imaging structure index, $\ln(1 + \text{PCI}) - \ln(1 + \text{coronary CT})$. The index falls when the front grows faster than treated output. H2 is therefore not only a ratio prediction. It is the operational claim that new upstream information does not immediately transmit one-for-one into downstream action.

Hypothesis 2 (H2). *After a hospital adopts FFR_{CT}, treatment output does not increase proportionately with upstream coronary CT imaging volume, so the treatment-to-imaging structure index decreases.*

Dynamic absorption (H3). The information-to-action relationship is dynamic. A care line has established referral norms, interpretation routines, scheduling practices, and cross-station coordination. A sudden increase in upstream information disrupts that fit. The front of the line can expand

immediately, while downstream actors need time to learn how to interpret, route, and incorporate the new information. This dynamic resembles the productivity J-curve for general-purpose technologies: local capability can improve before system output fully reflects complementary organizational adjustment (Brynjolfsson et al. 2021, Sterman et al. 2015).

In our setting, this logic predicts a J-curve in line structure. At adoption, the front expands faster than treatment output, so the treatment-to-imaging structure index falls. Over time, the line absorbs the new information-processing capability: referral rules, interpretation routines, and downstream conversion practices adjust. The prediction is about the shape of system adaptation, not about patient-level welfare. Public aggregate data can show that the line rebalances; they cannot show whether rebalancing reflects higher appropriateness or better clinical outcomes.

Hypothesis 3 (H3). *After a hospital adopts FFR_{CT}, the treatment-to-imaging structure index follows a J-curve: it decreases at adoption and recovers over subsequent years as the care line absorbs the new information-processing capability.*

Position as a boundary condition (H4). The first three hypotheses describe the effect of an AI placed at the upstream diagnostic gate. The theory is broader: the system effect should depend on position. If AI is placed at the front, it serves admission-relevant uncertainty and should mainly expand information production. If AI is placed closer to the downstream action margin, it serves a different uncertainty and should not generate the same front-end flood and diagnostic tilt.

The coronary line allows a within-line contrast. Physiological assessment is measured upstream through CT-derived fractional flow reserve and downstream through angiographic FFR in the cath lab. The contrast should be read as a boundary test rather than as the sole pillar of the paper, because the downstream AI arm is small. Its role is to ask whether the signature observed for FFR_{CT} is consistent with position rather than merely with the presence of AI.

Hypothesis 4 (H4). *The system effect of physiological AI differs by process position: upstream placement produces greater front-end expansion and a larger diagnostic tilt than downstream placement.*

3.4 The Binding Margin

The hypotheses imply a further operating question: which downstream margin governs output after upstream AI adoption? A capacity interpretation points to the cath lab or treatment stage. If treatment-stage capacity is binding, then additional downstream resources should unlock the front-end surge. A conversion interpretation points instead to a rate: the fraction of downstream arrivals that become treatment. If the binding margin is a rate, adding physical capacity does little unless the information arriving downstream becomes more treatment-relevant or the downstream decision process changes.

We do not state this as an additional hypothesis. The four hypotheses define the main theory chain: information production, weak immediate transmission, dynamic absorption, and positional boundary. The funnel model in Section 6 is a mechanism analysis that locates the reduced-form margins behind that chain. Because public aggregate data do not observe patient-level appropriateness, lesion severity, or clinical outcomes, the funnel is used to locate operational transmission rather than to judge whether the conversion rate is clinically optimal.

3.5 Summary of the Argument

The argument is a single information-processing mechanism. Diagnostic AI is a positioned capability. When placed at the gate, it increases information production at the front of the line (H1).

Because the downstream action margin is separate, the added information does not immediately produce proportional treatment output (H2). Because the care line must adapt its routines to the new information flow, the diagnostic tilt weakens over time and produces a J-curve (H3). Because the effect depends on which uncertainty the AI serves, changing the position of physiological AI should change the system signature (H4). The empirical sections test these four implications and then use the funnel to locate the margins through which the front-end shock is transmitted.

4. DATA AND EMPIRICAL STRATEGY

4.1 Data and the Adoption Signal

The signal we need is hard to see in public data, and most of the work in this section is in recovering it. All of our data are public Medicare files. We use no data-use agreement, no commercial source, and no patient-level record. The adoption signal and the flow outcomes both come from the Medicare Physician and Other Practitioners Public Use File ([Centers for Medicare & Medicaid Services 2024b](#)), which reports services at the level of the provider, the HCPCS code, and the year. This file is the only public source that resolves AI adoption because it records the specific HCPCS codes that the AI services carry, FFR_{CT} under HCPCS 75580 and the coronary CT it reads under 75572 through 75574, of which 75574, the coronary CTA code that also carries pre-procedural coronary imaging, is the large majority (about 77% of these services, with a share stable across 2016–2024 as cardiac CT for valvular and structural indications grew), while 75572 (cardiac morphology) and 75573 (congenital) are the smaller non-coronary remainder (about 23%) that the hospital fixed effects absorb in level; restricting the count to 75574 alone leaves the front-end flood and the structure tilt intact and if anything larger (Appendix B.8). The outpatient alternative does not: [Centers for Medicare & Medicaid Services \(2024a\)](#) documents that the CMS Outpatient by Provider and Service file bundles services into Ambulatory Payment Classification groups rather than reporting HCPCS codes, so it cannot isolate a single AI code. We checked this against our own extraction rather than relying on the documentation. The AI codes disappear once outpatient services are grouped into APC bundles, which leaves the Physician file as the only public source that identifies adoption.

Three further public files supply structure and controls. FFR_{CT} is read by the cardiologists and radiologists who staff a hospital's coronary line, and public data attach that measurement to the clinician rather than to the hospital. The Medicare Facility-Affiliation file records where each clinician practices, and we use it to read a hospital's coronary line through the AI activity of the clinicians who staff it. Because a clinician who staffs several hospitals is credited to each, this attribution smears capability and volume across sites and so attenuates adoption timing; the fractional-attribution panel (inverse affiliation count) brackets that choice and the random-timing placebo shows the smearing does not manufacture the result. The Healthcare Cost Report Information System (HCRIS) supplies time-varying hospital controls such as size and teaching status. The archived Care Compare files supply hospital descriptors used in sample construction. None of these files contain a patient record.

The provider file carries one coding rule that shapes what we can see, and handling it correctly is what makes the adoption signal trustworthy. Medicare suppresses any provider-code-year cell with fewer than eleven services. The rule removes the smallest billing counts, which truncates the adoption signal at small hospitals where a new service may bill only a handful of times in its first year. Aggregating provider billing to the hospital through the Facility-Affiliation link recovers part of the lost signal, because a count suppressed under one provider can clear the threshold once the affiliated providers are summed to the hospital. The aggregation cannot manufacture an observation below the threshold. It can only reassemble counts that suppression split apart. The

hospital is therefore the unit at which we measure flow, structure, and adoption.

We build the hospital-year panel (2016–2024) under three clinician-to-hospital attribution rules. The capability panel is the workhorse (4,148 hospitals, 35,835 hospital-years); the two-arm panel (2,423 hospitals) frames the upstream-versus-downstream position contrast; the balanced sub-panel of 3,769 hospitals (33,921 hospital-years) serves the synthetic difference-in-differences estimator; and the full-line core panel of 34,228 hospital-years, the subset of the capability panel on which the front, the invasive arrivals, and the output are jointly observed, serves the funnel of Section 6. The output is the percutaneous-intervention count; that this public count captures treated output rather than a billing artifact is established in Section 6 through the percutaneous-minus-bypass contrast and bounded in the Limitations, where the public-count limits run against the headline rather than toward it. In the capability panel, 1,201 hospitals adopt FFR_{CT} between 2018 and 2024 and form the treatment spine, while the standalone position-AI cells are thin: 74 hospitals carry downstream angiographic-FFR AI, of which only 23 have a post-adoption year inside the window, which is the boundary we keep in view when reading the downstream coefficients. The full construction, the Medicare small-cell suppression and how aggregation to the hospital recovers the signal, and the complete sample accounting (Table A1) are in Appendix A.

4.2 The Estimand and the Estimating Equation

We measure what AI at the gate does to the coronary care line on a hospital-year panel that runs from 2016 to 2024. The treatment is the staggered adoption of FFR_{CT}, the software that computes the fractional flow reserve of a coronary CT that a hospital has already acquired. Of the 4,148 hospitals in the capability-attribution panel, 1,201 take up FFR_{CT} at some point between 2018 and 2024, and the rest never adopt and remain as controls. The estimation is an operating-system effect, not a clinical one. We ask how the gate AI changes the line's volume and the balance between its diagnostic and treatment ends, and how both move over time, taking the tool's physiological accuracy as established by the validation trials cited in Section 2.2 rather than re-evaluating it. Two designs could speak here, and they pull apart cleanly. A trial randomly selects patients and measures their health. We observe hospitals adopting the tool in the field and measuring what their line does, and it is that operating effect that we identify.

The baseline specification is a two-way fixed-effects regression with hospital and year effects,

$$Y_{ht} = \alpha_h + \lambda_t + \beta \text{Post}_{ht} + \delta_U \text{Post}_{ht} \times \text{Up}_h + \delta_D \text{Post}_{ht} \times \text{Down}_h + X_{ht}\gamma + \varepsilon_{ht}, \quad (1)$$

where Y_{ht} is the outcome for hospital h in year t ; α_h and λ_t are fixed effects for the hospital and year; Post_{ht} turns on in the FFR_{CT} adoption year and stays on; Up_h and Down_h indicate upstream (plaque) and downstream (angiographic-FFR) AI capability; X_{ht} collects time-varying HCRIS controls; and standard errors are clustered at the hospital. We read three outcomes in the order the line runs. The flow index (H1) is $\ln(1 + \text{coronary CT})$, the volume of admission-relevant sensing the gate performs. The structure index, $\ln(1 + \text{PCI}) - \ln(1 + \text{coronary CT})$, is the balance between the treatment and diagnostic ends of the line, registering whether the information-processing capacity added at the gate reaches the requirement that binds at the converting station, and carries both the prediction of the structure (H2) and the prediction of the path of time (H3). The output index is $\ln(1 + \text{PCI})$, the treatment the line completes once the converting station's detection capacity has acted, and the position contrast (H4) is read across these outcomes through the Up and Down interactions. We prefer the structure index to a raw conversion ratio because it treats the numerator and denominator symmetrically and does not mechanically reward a hospital whose coronary-CT count increases. The design is a difference-in-differences, and everything that follows turns on which counterfactual the data can support; that support is strong for some outcomes and thin for

others.

4.3 Selection into Adoption

The threat to this design is that hospitals do not adopt FFR_{CT} at random and the way they select is tied to the front of the line by construction. FFR_{CT} is post-processing software that runs on a coronary CT already in hand, so a hospital cannot adopt it without a coronary CT program already in place; the tool analyzes records of previously acquired computed-tomography data ([National Institute for Health and Care Excellence 2017](#)). [David \(1969\)](#) shows, for a technology that carries a fixed adoption cost, that the adoption occurs once the relevant scale crosses the threshold at which the cost is justified, and [Karshenas and Stoneman \(1993\)](#) formalize this as the rank mechanism of diffusion, in which a unit adopts when its own returns cross the cost margin. Hospital imaging adoption follows that pattern, increasing with program volume and size, so a hospital adopts the gate AI as its front-end volume is already climbing.

An option-value adoption-trigger model formalizes this timing (Appendix A.4): a hospital adopts once its front-end volume first crosses an investment trigger, so adopters' pre-adoption front-end trajectory slopes upward by construction ([Dixit and Pindyck 1994](#), [Ashenfelter 1978](#), [Heckman et al. 1999](#)). The observed jump in front-end volume after adoption could therefore come from the very trend that triggered adoption, not from the tool.

This selection is visible in the leads. A joint test of the pre-adoption leads rejects parallel trends on both level outcomes ($p < 0.001$), and in the structure index the lead one year before adoption (-0.354) is already larger in magnitude than the adoption-year effect (-0.309), so the dip at adoption continues a decline that is underway before the tool arrives. This is exactly what the adoption-trigger mechanism predicts: hospitals take up FFR_{CT} as their front-end volume is rising, so a non-parallel level pre-trend is the expected pattern, not an anomaly. We therefore identify the level effects not from the raw two-way fixed-effects event study of equation (1) but from estimators built for selection on the outcome's own trajectory (Section 4.4).

The same leads point to the design that does identify cleanly. While the level leads are not parallel, the position-moderating interactions are: the joint test on the downstream-AI increment leads gives $p = 0.539$ and on the upstream-AI increment leads $p = 0.666$, so the pre-adoption dynamics do not differ between position-distinct adopters. The downstream test is low-powered given the thin arm, so we read its non-rejection as no detectable pre-trend rather than a demonstrated flat one. The selection the adopters share therefore cancels in the contrast even though it moves the levels, and that contrast is the design the rest of the paper builds on (Section 4.5).

4.4 The Estimator Ladder

Because the selection is anticipated, we identify the effects from estimators built for settings where adoption tracks the outcome's own pre-period path, and read them as a ladder rather than as a single coefficient. Synthetic difference-in-differences ([Arkhangelsky et al. 2021](#)) on the balanced sub-panel supplies the level effects: it raises upstream flow by $+0.357$ log points, lowers the structure index by -0.299 , and raises treated output $\ln(1 + \text{PCI})$ by $+0.074$ (bootstrap 95% CI $[0.035, 0.114]$), all significant ($p < 0.001$; Table 1). The [Borusyak et al. \(2024\)](#) imputation estimator traces the event-time path of the structure effect (Section 5.3). The [Callaway and Sant'Anna \(2021\)](#) not-yet-treated estimator supplies the selection-robust member, at $+0.204$ for flow, -0.131 for structure, and $+0.073$ for output (all $p < 0.001$); because it builds the counterfactual only from hospitals not yet adopting, it sits at the lower-magnitude end of the selection bracket that [Ding and Li \(2019\)](#) characterize. The adoption-trigger mechanism, a hospital adopts when its own front-end returns cross a fixed cost threshold ([David 1969](#), [Karshenas and Stoneman 1993](#)), is a monotone-index

rule, the monotone-selection condition under which [Ding and Li \(2019\)](#) place this lower-magnitude member at the end of the bracket whose sign agreement with the preferred estimate we read rather than a guaranteed floor. The [Anderson and Hsiao \(1982\)](#) first-difference instrumental variable, at $+0.201$ for flow and -0.178 for structure, is a lower-magnitude member reported in Appendix B with its weak-instrument and serial-correlation diagnostics. Across the fuller ladder, which adds two-way fixed effects, the [Sun and Abraham \(2021\)](#) staggered estimator, fractional attribution, and hospital-specific trends, the flow effect stays inside $[+0.20, +0.63]$ and the structure effect inside $[-0.57, -0.18]$, every member keeping its sign and clearing zero (Table B1, Figure B1). Table B2 crosswalks each headline coefficient to the panel and estimator that identifies it and the parallel-trends or sensitivity diagnostic computed on that pairing. The estimator mechanics are in Appendix B.

4.5 Identification of the Position Contrast

The position contrast is recovered more cleanly than the levels, and it is what carries the claim that placement, not the mere presence of AI, governs the system effect. The selection into FFR_{CT} adoption is common to the adopters, so it differences out of a comparison between adopter types, and the interaction leads show no detectable pre-trend: the joint tests on the position increments do not reject (Section 4.3), with the caveat that the downstream arm is thin, so its non-rejection reflects limited power as much as a literally flat lead. The contrast therefore does not inherit the level pre-trend, and the separation is identified off the well-populated upstream arm; we lean on the sign reversal between positions and the upstream-powered separation rather than on a precise downstream pre-trend or a measured downstream null. This is the differencing-out logic set out by [Roth et al. \(2023\)](#): when a main effect identifies off contaminated leads, a moderation contrast can still identify off clean interaction leads, because the contamination the two arms share is removed by the difference between them. We report the contrast itself, the upstream-versus-downstream signature that is the climax of the empirical argument, in Section 5.4; here the data support the contrast.

4.6 Remaining Threats

Three threats could still qualify the identification, and the first we meet with direct evidence. The first is a shock that lands together with adoption: FFR_{CT} is software bolted to an installed coronary-CT scanner, so a cardiology program expanding its line could install the tool at the same moment, and a level effect could then reflect the expansion rather than the tool. Three pieces of evidence rule this out, all pointing against a coincident confound. The structure tilt deepens monotonically with cumulative FFR_{CT} dose, a within-adopter gradient a shock common to all adopters could not produce (Appendix B). A placebo built on the non-AI version of the same measurement runs the other way: among 2,298 hospitals that adopted wire-FFR but never installed FFR_{CT} , wire-FFR, an invasive in-cath-lab measurement billed by invasive-end-heavy programs, raises the structure index by $+0.694$ ($p < 0.001$), the opposite sign that its downstream position predicts; generic cardiology investment of this kind moves the line the other way, so it looks different from the gate AI rather than being absent, and it is one of three converging checks beside the within-adopter dose gradient and the ladder sign-robustness. And a [Rambachan and Roth \(2023\)](#) relative-magnitude sensitivity, once each hospital's adoption trajectory is absorbed, places the breakdown above the $\bar{M} = 1$ benchmark for both outcomes, so a coincident shock larger than the worst remaining pre-trend would be needed to overturn the level estimate (Figure C1; Appendix C). Because this relative-magnitude breakdown is computed on the two-way-fixed-effects event-study leads rather than on the synthetic estimator's weighted leads, it bounds the event-study estimand rather than the syn-

thetic level itself; we therefore certify the structure result through its sign-robustness across the full ladder ($[-0.57, -0.18]$, lower-magnitude first-difference IV member -0.178), the within-adopter dose gradient, and the wire-FFR opposite-sign placebo, not through the precise -0.299 alone (Section 5.2).

The second threat is sample coverage. The balanced sub-panel comprises larger and more stable hospitals and does not reach the small-hospital tail, so the level effects are identified for the institutions that adopt early and survive in the panel, and we read them as such. The third is the boundary of the data itself. We observe the volume and the balance of AI use along the diagnostic pathway, and the panel carries no patient-outcome measure, so the welfare of patients exposed to FFR_{CT}-induced upstream volume sits outside what we can identify and outside every claim we make. None of the three runs in the direction of our finding. The dose gradient, the flat anticipation contrast, and the wire-FFR reversal all cut against a coincident shock; the coverage limit narrows the population rather than the sign; and the outcome boundary caps the claim rather than inflating it.

5. RESULTS

We organize the results to match the theory chain. The first result shows that the upstream AI releases information-processing capacity at the diagnostic gate. The second shows that this information production does not immediately become proportional downstream action. The third shows that the initial diagnostic tilt is dynamically absorbed. The fourth asks whether the same system signature depends on AI position. This order matters: the position contrast is a boundary test; the main evidence for care-line absorption is the H1–H3 sequence.

Table 1: Main results: the system effect of the gate AI

	Flow index	Structure index	Output index	Arrivals index
Synthetic DiD (preferred)	+0.357*** (0.034)	−0.299*** (0.038)	+0.074*** [0.035, 0.114]	+0.029
Callaway–Sant’Anna (selection-robust)	+0.204***	−0.131***	+0.073***	+0.032**
Hospital and year effects	Yes	Yes	Yes	Yes
Proportional effect (preferred)	+43%	−26%	+8%	+3%

Notes. Preferred estimator: synthetic difference-in-differences (Arkhangelsky et al. 2021) on the balanced sub-panel (33,921 hospital-years, 3,769 hospitals), with bootstrap inference. Selection-robust check: Callaway and Sant’Anna (2021) not-yet-treated estimator on the capability panel (34,228 hospital-years). SE in parentheses; 95% CI in brackets for output. Proportional effects are $\exp(\hat{\beta}) - 1$; structure implies $\exp(-0.299) - 1 = -25.8\%$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.1 Front-End Information Production (H1)

The first effect is local to the station where the AI is deployed. In the balanced sub-panel, the synthetic difference-in-differences estimate for upstream flow is +0.357 log points (SE 0.034, $t = 10.4$, $p < 0.001$, 95% CI [0.290, 0.424]). In proportional terms, adoption increases front-end coronary CT activity by about 43%. The estimate is substantively large and precisely estimated, and it is the empirical counterpart of H1: the gate AI expands information production at the front of the care line.

The result is not specific to one estimator. Across the estimator ladder in Section 4, the flow effect remains positive, significant, and within the bracket [+0.20, +0.63]. The conservative first-

difference instrumental-variable estimate is $+0.201$ ($p < 0.001$). Estimator choice changes the magnitude, but not the conclusion. The new information-processing capability is first observed where the theory predicts it should appear: at the diagnostic gate.

5.2 Weak Immediate Information-to-Action Transmission (H2)

The front-end expansion is not matched by proportional downstream action. The synthetic difference-in-differences estimate for the structure index, $\ln(1+PCI) - \ln(1+\text{coronary CT})$, is -0.299 log points (SE 0.038, $t = -7.8$, $p < 0.001$, 95% CI $[-0.375, -0.224]$). Treatment per unit of imaging falls by about one quarter after adoption. The line therefore becomes more diagnostic-intensive at adoption.

The structure index is mechanically related to front-end and output levels, so we do not treat it as a separate third outcome. Its theoretical content is the transmission failure it summarizes. The front rises by $+0.357$ log points, while output rises by only $+0.074$ log points in the preferred specification and by roughly $+0.03$ to $+0.07$ across selection-robust estimators. The front-end information shock therefore reaches downstream action only weakly. Across the estimator ladder, the structure estimate remains negative in the bracket $[-0.57, -0.18]$, including the conservative first-difference IV member at -0.178 ($p < 0.001$). This supports H2: upstream information production expands faster than downstream action.

5.3 Dynamic Absorption: The J-Curve (H3)


The diagnostic tilt is not a permanent step down. The event-time path shows a J-curve in line structure. Under the interaction-weighted estimator of [Sun and Abraham \(2021\)](#), the structure index falls to -0.321 in the adoption year ($p < 0.001$), remains below the counterfactual after one and two years (-0.181 , $p < 0.001$; -0.118 , $p = 0.019$), and is no longer distinguishable from the counterfactual by the third and fourth years (-0.049 , $p = 0.453$; -0.005 , $p = 0.946$). Figure 3 plots this path.

The interpretation is line-level absorption. The front expands immediately, downstream action lags, and the balance of the line then recovers over time. A random-timing placebo that reassigns adoption years among treated hospitals does not reproduce the observed dip-and-recovery path in 200 draws ($p < 0.005$). The recovery also survives excluding the 2020–2021 pandemic years and absorbing differential pandemic shocks through pre-adoption-volume-tertile-by-year fixed effects (Appendix B.4). The late event-time points are imprecise, and the public data cannot show patient-level improvements in appropriateness or safety. The claim is therefore operational: the aggregate care-line structure absorbs the new upstream information-processing capability over the observed window.

The cross-sectional pattern is consistent with this dynamic reading. Treated-output responses do not concentrate in hospitals with the resources a simple capacity-bottleneck account would emphasize—bed size, downstream catheterization capacity, teaching status, surgical capability, ownership, safety-net status, or local disease burden. The limited heterogeneity that appears tracks region, market adoption, and hospital quality rather than treatment-stage resources (Appendix D). This evidence does not prove a patient-level mechanism, but it supports the view that the J-curve reflects care-line adjustment rather than a fixed physical bottleneck.

5.4 Position as a Boundary Test (H4)

The first three results establish the main mechanism for upstream AI. H4 asks whether that signature is positional. In the two-arm position contrast, upstream AI produces a large front-end effect and a diagnostic tilt: coronary imaging rises by $+0.463$ log points, while the structure index falls



Notes. [Borusyak et al. \(2024\)](#) imputation estimate of the structure index by event time (navy, 95% confidence intervals); [Sun and Abraham \(2021\)](#) estimator (blue dashed); decaying-selection projection of the dip (gray dashed), the path the structure index would trace if the adoption-year dip were pure selection decaying geometrically toward the counterfactual. Event time 0 is the adoption year.

Figure 3: Event-time path of the structure index.


by -0.404 . Downstream AI does not produce the same signature: its flow effect is -0.134 and its structure effect is $+0.223$, neither statistically distinguishable from zero in the thin downstream arm (Table 2; Figure 4). The two positions separate at $p < 0.001$ for flow and structure.

Table 2: Two-arm position contrast.

	(1) Flow index	(2) Arrivals index	(3) Output index	(4) Structure index
Upstream AI (plaque), post	0.463*** (0.0593)	0.0113 (0.0247)	0.0588 (0.0383)	-0.404*** (0.0682)
Downstream AI (angio-FFR), post	-0.134 (0.148)	0.0543 (0.0692)	0.0890 (0.128)	0.223 (0.163)
Observations	15152	15152	15152	15152
Within R-squared	0.824	0.931	0.865	0.784

Notes. Hospital and year fixed effects; SE clustered by hospital in parentheses. Indices: flow = $\ln(1 + \text{coronary CT})$, arrivals = $\ln(1 + \text{invasive coronary angiography})$, output = $\ln(1 + \text{PCI})$, structure = $\ln(1 + \text{PCI}) - \ln(1 + \text{coronary CT})$. Upstream arm: plaque AI (0623T–0626T); downstream arm: angiographic FFR (0523T). Wald test of equality rejects at $p < 0.001$ for flow and structure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The downstream arm contains 74 hospitals, with only 23 contributing a post-adoption year. We therefore read the downstream estimates as bounding the downstream effect rather than precisely estimating a zero. The inference rests on the sign reversal and the between-position separation, which are powered by the well-populated upstream arm and supported by clean interaction leads.



Notes. Two-arm panel; hospital and year fixed effects, SE clustered by hospital. Points are post-adoption effects with 95% confidence intervals; filled markers are significant at $p < 0.05$. The flow index is $\ln(1 + \text{coronary CT})$ and the structure index is $\ln(1 + \text{PCI}) - \ln(1 + \text{coronary CT})$. The downstream arm is thin (74 hospitals, 23 with a post-adoption year); at 80% power it detects only effects beyond about 0.41 log points, so it bounds the downstream effect rather than measuring a zero, and the claim is the sign reversal and the between-arm separation ($p < 0.001$).

Figure 4: Position contrast by AI placement.

A better-powered within-adopter interaction, estimated among FFR_{CT} adopters with downstream-AI capability, yields the same positive sign on the structure index and is reported with leave-one-out checks in Appendix B.

The position contrast should therefore be read as a boundary test. It does not carry the entire paper. The main evidence is the H1–H3 sequence showing front-end information production, weak immediate transmission, and dynamic absorption. H4 adds that the same care-line signature is not a generic consequence of AI adoption; it depends on where the information-processing capability is embedded.

5.5 Heterogeneity by Line Shape

The average front-end expansion masks different responses by pre-existing line shape. In cath-forward lines, where the invasive end is already large relative to the diagnostic front, the gate AI is associated with substitution: invasive arrivals and treatment contract ($\ln(1 + \text{ICA})$, -0.051 , $p = 0.004$; $\ln(1 + \text{PCI})$, -0.051 , $p = 0.040$). In CT-forward lines, the same AI expands the line and pulls downstream activity with the enlarged front. These regimes support the positional argument by showing that the system effect is not determined by the tool alone. It depends on the structure of the line the tool enters. Full regime definitions and estimates are in Appendix E.

5.6 Motivation for the Funnel Mechanism

The results establish a system signature but not the margin that produces it. Front-end activity expands, downstream action responds weakly, the line structure recovers over time, and the pattern differs by AI position. The remaining question is where the front-end shock is lost. A manager might look to the cath lab, the scarce and scheduled station near the end of the line. The reduced form also allows other explanations: the front-end surge may be throttled before invasive arrival, or it may reach the invasive stage but fail to convert into treatment. Section 6 reads the line as a funnel to locate these margins within the public aggregate data.

6. MECHANISM: A FUNNEL MODEL OF THE CARE LINE

The results identify the shape of the operating effect. The funnel locates where the front-end shock is attenuated. We use the funnel as a reduced-form decomposition, not as a structural model and not as patient-level evidence. Its purpose is to separate two margins that are observationally blended in the headline estimates: pass-through from the diagnostic front to the invasive stage, and conversion from invasive arrival to treatment.

A conventional capacity account points to the treatment end of the line. If the cath lab or treatment stage is binding as a capacity level, then expanding downstream capacity should release output from the enlarged diagnostic front. The information-processing account points to a different object. If output is governed by a conversion rate, then more front-end candidates and more physical capacity will not raise output unless the downstream action margin converts a larger fraction of arrivals. The funnel adjudicates this level-versus-rate distinction in the public aggregate data.

6.1 A Reduced-Form Funnel

The line carries flow through three observed stages. Let f denote front-end coronary CT volume, A invasive angiography arrivals, and y treated output, measured primarily as PCI. The first margin is pass-through from the front to the invasive stage. The second is conversion from invasive arrival to treatment. We write the reduced-form funnel as

$$A = D + \tau f, \quad (2)$$

and

$$y = \rho A, \quad (3)$$


where D is baseline referral activity, τ is the pass-through margin, and ρ is the descriptive conversion rate. In log specifications we estimate a conversion elasticity η , the percentage change in output associated with a percentage change in arrivals. A gate AI reaches output only through the product of these two margins: the front-end shock must pass through to invasive arrivals, and those arrivals must convert into treatment.

The model is deliberately parsimonious. It uses ratios of reduced-form coefficients to locate operational transmission. It does not recover a structural patient-level transition probability and it does not judge whether conversion is clinically appropriate. It asks the operational question that the results leave open: how much of the front-end shock survives to each downstream stage?

6.2 Estimating the Two Margins

The estimated pass-through is small. In the balanced panel, synthetic difference-in-differences estimates imply a front-end increase of +0.357 log points, an invasive-arrivals increase of +0.029, and an output increase of +0.074. The implied front-to-arrival pass-through is $0.029/0.357 = 0.08$, about eight percent of the front-end shock. In the full-line core panel, two-way fixed effects produce a larger front estimate (+0.588) and a similar arrivals estimate (+0.033), implying pass-through of $0.033/0.588 = 0.06$ (delta-method 95% CI [0.01, 0.10]). Across these readings, only about six to eight percent of the diagnostic surge reaches the invasive stage.

Treated output rises, but modestly. The synthetic estimate is +0.074 log points, the core-panel fixed-effects estimate is +0.069, and selection-robust estimates bracket the output response at roughly +0.03 to +0.07. Thus, output is not zero; it is small relative to the front-end increase. This distinction is important for interpretation. The paper does not claim that upstream AI produces no downstream action. It claims that action transmission is weak relative to the information-production shock.



Notes. Schematic. τ is the front-to-cath-lab pass-through and ρ the conversion rate beyond it; estimated magnitudes are in Section 6.

Figure 5: The care line as a funnel (schematic).

The conversion elasticity is close to proportional. The estimate of η is 0.954 ($p < 0.001$), meaning that a one percent increase in invasive arrivals is associated with about a one percent increase in output. The adoption interaction is +0.045 (SE 0.035), and we cannot reject zero ($p = 0.196$). An equivalence test against a 10% band rejects shifts beyond that band at the ten-percent level ($p = 0.075$). We therefore treat the conversion margin as near-invariant to adoption in the counterfactuals, while carrying the uncertainty in sensitivity checks. Table 3 reports the estimates.

Table 3: Conversion frontier and the transmission cascade

	(1) Front	(2) Arrivals	(3) Output	(4) Conversion
Gate AI adoption	0.588*** (0.038)	0.033** (0.013)	0.069*** (0.017)	
Gate AI adoption (SDiD)	0.357*** (0.034)	0.029	0.074*** [0.035, 0.114]	
Conversion elasticity η				0.954*** (0.031)
Gate AI (intercept shift)				-0.251 (0.232)
η shift at adoption				0.045 (0.035)
Observations	34228	34228	34228	34228
AI-invariance test p				0.196

Notes. Row 1: two-way fixed effects on the core panel (34,228 hospital-years). Row 2: SDiD on the balanced panel (33,921 hospital-years) with bootstrap inference. SE in parentheses; 95% CI in brackets for output. Pass-through: 0.08 (SDiD) and 0.06 (FE). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These quantities clarify the J-curve. At adoption, the front expands sharply, pass-through to invasive arrivals is small, and the conversion margin does not detectably shift. Over time, the structure index recovers. The static funnel cannot itself generate the dynamic path, but it locates the margins through which adjustment must occur: pass-through can normalize, conversion can improve, or both. The J-curve is therefore the dynamic counterpart of the funnel. It shows that the information-to-action relationship is absorbed gradually rather than instantaneously.

6.3 The Binding Margin

The capacity interpretation receives limited support. If treatment-stage capacity were the binding level, output should flatten where downstream load or hospital capacity is high. Instead, output rises roughly in step with invasive load in the bed-data subset (elasticity 0.926, $p < 0.001$), and the

load-output relation does not show the saturation pattern expected of a binding treatment-stage capacity. The same conclusion holds when using cath-lab procedural capacity, hospital-specific trends, and downstream capacity moderators (Appendix F; Appendix D). These are indirect tests because public data do not observe cath-lab occupancy or slack, but they do not point to treatment-stage capacity as the margin that unlocks the front-end surge.

The reduced-form ranking points instead to transmission and conversion rates. The gate AI changes f , but the pass-through τ is small and the conversion margin is not detectably moved at adoption. Under the core-panel estimates, the pass-through-then-convert channel carries about six percent of the front-end shock. The raw output-to-front ratio is larger, about twelve percent, because the one-channel projection does not capture every level adjustment in the line. Under either reading, the conclusion is the same: most of the diagnostic expansion is not translated into downstream action within the observed window.

This is the operational meaning of a rate-like binding margin. Adding more front-end diagnostic volume does not move output proportionately because the line discards most of the front-end shock before treatment. Adding downstream physical capacity would also be limited if the downstream stage is not binding as a level. The lever that would change output lies in the information-to-action margins: referral pass-through and downstream conversion. A gate AI may support screening and deferral, but it does not automatically change those margins at adoption.

The public data limit the interpretation. We cannot observe whether the conversion rate is clinically correct, whether deferred patients were appropriately managed, or whether the downstream patient pool became more treatment-relevant. The funnel therefore locates operational transmission rather than welfare. It explains why the aggregate line exhibits front-end expansion, weak immediate transmission, and later re-equilibration.

7. ROBUSTNESS

The empirical challenge is that hospitals do not adopt FFR_{CT} randomly. Adoption is tied to the front of the line because a hospital must already have coronary CT activity before the AI can be used. The robustness analysis therefore focuses on whether the H1–H3 pattern is an artifact of estimator choice, adoption timing, generic cardiology investment, pandemic disruption, physician attribution, or influential downstream-AI hospitals.

The conclusions are stable. Across the estimator family, FFR_{CT} adoption increases front-end coronary CT volume and lowers the pathway-structure index. The signs survive synthetic difference-in-differences, not-yet-treated staggered-adoption estimators, first-difference IV, fractional attribution, and hospital-specific trend specifications. A wire- FFR placebo moves the structure index in the opposite direction, suggesting that generic invasive cardiology investment does not explain the gate-AI signature. The early structure dip appears under both imputation and interaction-weighted staggered-adoption estimators, and no random-timing placebo reproduces the observed J-curve.

The position contrast is also not driven by one hospital. Leave-one-out checks preserve the upstream-versus-downstream separation, although the downstream arm remains thin and should be read as a boundary test. Dropping the 2020–2021 pandemic years, using fractional physician-to-hospital attribution, and imposing balanced pre-period restrictions leave the main signs and orders of magnitude intact. Taken together, the robustness checks support the public-data claim: upstream diagnostic AI expands information production, the expansion is weakly transmitted to downstream action at adoption, and the line structure absorbs the shock over time.

8. DISCUSSION

8.1 Contribution to Information-Processing Theory

The main theoretical contribution is to reposition AI within information-processing theory. Classical theory argues that organizations facing uncertainty need additional information-processing capacity or a better fit between information demand and information processing (Galbraith 1974, Tushman and Nadler 1978). AI seems to provide exactly that capacity. Our evidence shows why this statement is incomplete for sequential care lines. AI is not only an amount of information-processing capacity. It is a positioned capability. It serves the informational requirement of the station where it is embedded, and its system value depends on whether that station is close to the margin that converts information into action.

The coronary line makes the distinction visible. The diagnostic-gate AI expands information production at the front by about 43%. Yet only a small fraction of this increase reaches the invasive stage, and treated output rises only modestly. This pattern does not mean the tool is clinically ineffective; FFR_{CT} is designed partly to improve triage and avoid unnecessary invasive testing. It means that local diagnostic information and system output are governed by different margins. The front of the line can become information-rich while treatment remains governed by downstream action conversion.

The J-curve adds a dynamic extension to IPT. The system effect of AI does not arrive as a single new level. It arrives as a sequence: immediate front-end expansion, diagnostic tilt, and gradual rebalancing. This is line-level absorption. The organization must learn how to interpret the new information, route patients under the new signal, schedule downstream work, and incorporate the signal into treatment decisions. Public aggregate data cannot show patient-level learning or appropriateness, but they can show the operating signature of absorption: the structure index falls and then recovers.

8.2 Relation to Constraints Theory

A simple constraints interpretation would say that improving a non-bottleneck stage cannot improve system output. That is part of the result, but it is not the full contribution. The upstream diagnostic gate is not where treatment output is completed, so proportional throughput gains should not be expected. But the funnel shows something more specific: the relevant margin behaves more like a rate than a physical capacity level. Output is not limited primarily by observable treatment-stage capacity in our data; it is limited by the rate at which candidates pass through and convert downstream.

This difference matters for managerial prescription. If the bottleneck were a capacity level, a hospital could respond to front-end AI by expanding downstream slots. If the binding margin is a conversion rate, adding slots does not solve the transmission problem. The operating lever must change referral pass-through, downstream interpretation, or conversion rules. In information-processing terms, the AI should be aimed at the uncertainty that governs action, not simply at the earliest uncertainty in the line.

This also reframes the phrase “misplaced intelligence.” The intelligence is not misplaced because the algorithm is wrong or because the clinical tool lacks value. It is misplaced relative to a throughput objective when it is installed upstream of the margin that governs output. The same AI can still create value through screening, deferral, and re-sorting. The point is that those benefits appear as structural reallocation and delayed absorption, not as immediate treated throughput.

8.3 What the Gate AI Does

The gate AI does three things. First, it relocates work to the diagnostic front. The front-end increase is not saved effort; it is additional information production. This is consistent with the economics of cheaper sensing and with supply-induced demand in healthcare (Roemer 1961, Gruber and Owings 1996, Clemens and Gottlieb 2014, Agrawal et al. 2019).

Second, it changes the balance of the line. Treatment per unit of imaging falls at adoption because the front expands faster than output. This diagnostic tilt should not be read mechanically as failure. In the clinical logic of FFR_{CT}, part of the value of the tool is to avoid unnecessary invasive testing. In our data, the claim is narrower: aggregate line activity shifts toward diagnosis, and public counts cannot determine whether that shift improves appropriateness.

Third, the tilt is absorbed over time. The structure index re-equilibrates toward its counterfactual within the observed window. This recovery is the core dynamic result. It indicates that a care line does not passively receive AI-generated information; it reorganizes around it. The public data do not identify the precise organizational routine through which absorption occurs, but the time path is inconsistent with a purely static front-end shock.

8.4 Position as a Design Variable

The position contrast supports the broader design logic. Upstream AI produces a front-end expansion and diagnostic tilt. Downstream AI does not produce the same signature, although the downstream estimates are imprecise and should be interpreted as bounds. This evidence is sufficient for the role it plays in the paper: it shows that the system signature is consistent with process position, not merely with the presence of AI.

Position is therefore a design variable. Choosing where to deploy AI is choosing which uncertainty the system will process more cheaply. An upstream tool targets admission-relevant uncertainty and therefore changes candidate generation. A downstream tool targets action-relevant uncertainty and should have a different operating signature. This distinction complements clinician-level work on whether clinicians should rely on AI (Dai and Singh 2025). A tool can be useful at the local clinical decision and still have limited effect on system output if it is placed far from the downstream conversion margin.

Line shape further conditions this effect. In CT-forward lines, the front can feed downstream expansion. In cath-forward lines, the gate AI is associated with substitution at the invasive end. The same tool therefore enters different operating systems and produces different responses. This contingency is consistent with JOM healthcare evidence showing that utilization, focus, and mix uncertainty shape how hospitals absorb operating complexity (Thirumalai and Devaraj 2023, Thirumalai et al. 2024).


8.5 Managerial Implications

The main managerial implication is that diagnostic AI should be procured and evaluated as a care-line intervention, not only as a local diagnostic tool. If the goal is better triage, a front-end AI can be valuable even when treatment volume does not rise. If the goal is more treated output, a gate AI alone is unlikely to suffice unless the downstream action margin can also change.

Hospitals should therefore evaluate three quantities after adoption: how much front-end activity expands, how much of that activity reaches the next stage, and how the line structure evolves over time. First-year evaluation is especially risky. In our data, the first-year pattern looks like a diagnostic tilt, but the structure later re-equilibrates. A short evaluation horizon can misclassify a dynamic absorption process as either failure or success.

Payers and guideline writers face a parallel issue. Reimbursing a scalable diagnostic gate

funds information production at the front of the line. That may be appropriate if the goal is safe deferral and better sorting, but it should not be confused with buying downstream output. Payment and guideline evaluation should therefore ask which margin the tool is expected to move: diagnostic information production, referral pass-through, treatment conversion, or patient outcomes.



Notes. Input–output space indexed to the pre-adoption point (= 100, 100); a ray from the origin is treated output per unit of imaging, so a steeper ray is more efficient. The filled point marks the estimated post-adoption position; the dashed arrow marks the conversion-margin direction.

Figure 6: Operating efficiency of the care line.

Figure 6 summarizes the operating implication. Gate AI moves the line mostly rightward: more front-end imaging with a smaller movement in output. A tool or complementary redesign that moves the conversion margin would shift the line upward. The managerial lesson is not that the front-end tool lacks value. The lesson is that value appears at the margin the tool can reach.

8.6 Generalization Beyond Coronary Care

The positional mechanism should travel beyond coronary care wherever scalable front-end judgment feeds a downstream action step. Radiology reads precede biopsies, screenings precede interventions, AI-scored applications precede approvals, and AI-qualified leads precede sales. In each case, upstream AI can expand candidates without proportionally increasing completed action if the downstream conversion margin remains unchanged.

The theory is also falsifiable. If downstream conversion moves with the front—for example, if biopsies rise in step with AI-driven imaging or approvals rise in step with AI-scored applications—then the action margin is moving with information production and the positional mechanism is not binding. The boundary condition is therefore clear: the mechanism should hold when downstream action is governed by a relatively fixed rate or by routines that require absorption; it should weaken when the downstream decision itself is automated or highly scalable.

8.7 Limitations

Three limitations bound the claims. First, the data identify operating structure, not patient welfare. We observe aggregate hospital-year volumes of imaging, invasive evaluation, and treatment. We do not observe patient-level appropriateness, lesion severity, symptoms, or adverse outcomes. Thus, the paper cannot determine whether induced imaging is clinically desirable, whether deferred patients were appropriately managed, or whether the recovered line structure reflects better patient selection. External clinical evidence supports the safety and triage value of CT-first pathways, but our own evidence is operational rather than clinical ([The SCOT-HEART Investigators 2018](#), [Fairbairn et al. 2018](#), [Zink et al. 2026](#)).

Second, adoption is not random. Hospitals adopt FFR_{CT} when the front of their line is already growing, so level estimates rely on methods designed for selection on pre-period trajectories. We therefore present estimator brackets and robustness checks rather than treating any single coefficient as definitive. The dynamic J-curve and position contrast are central because they rely on patterns that are harder to explain by a simple adoption trigger alone.

Third, public data impose measurement limits. Physician-to-hospital attribution can smear activity across facilities, small-cell suppression can truncate low-volume services, and annual aggregation hides short-run dynamics. The downstream AI arm is thin, so H4 is a boundary test rather than the empirical foundation of the paper. Longer horizons and patient-level claims would allow tests of clinical appropriateness and safety, but they are not necessary for the paper's main public-data claim: upstream AI changes line structure through information production, weak immediate transmission, and dynamic absorption.

9. CONCLUSION

This paper examines diagnostic AI as a positioned information-processing capability in a sequential care line. In coronary care, FFR_{CT} adoption substantially increases front-end imaging, but only a small share of that increase reaches invasive evaluation and treated output rises only modestly. The line initially becomes more diagnostic-intensive and then re-equilibrates over time. The pattern is a J-curve in care-line structure: information production expands first, downstream action transmission lags, and the line absorbs the new capability gradually.

The main implication is theoretical and operational. AI's value is not determined only by local algorithmic performance. It depends on where the information is produced relative to the margin that turns information into action. When AI is placed at the diagnostic gate, its most visible operating effect may be diagnostic abundance and structural rebalancing rather than immediate throughput. When the goal is treated output, managers must identify and move the downstream action margin, not only automate the upstream gate.

The paper does not claim that weak immediate downstream transmission is clinically good or bad. Public aggregate data cannot answer that question. The contribution is instead to show how a validated diagnostic AI reshapes the operating line into which it is inserted. For healthcare organizations adopting AI at scale, this distinction is central: local intelligence becomes system value only when the operating system can absorb and act on the information it produces.

REFERENCES

- | | |
|--|---|
| <p>Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. <i>Journal of Economic Perspectives</i>, 33(2):31–50.</p> | <p>Anderson, T. and Hsiao, C. (1982). Formulation and estimation of dynamic models using panel data. <i>Journal of Econometrics</i>, 18(1):47–82.</p> <p>Arkhangelsky, D., Athey, S., Hirshberg, D. A., Im-</p> |
|--|---|

- bens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.
- Ashenfelter, O. (1978). Estimating the effect of training programs on earnings. *The Review of Economics and Statistics*, 60(1):47.
- Benjamins, S., Dhunnoo, P., and Meskó, B. (2020). The state of artificial intelligence-based fda-approved medical devices and algorithms: an online database. *npj Digital Medicine*, 3(1).
- Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.
- Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Centers for Medicare & Medicaid Services (2024a). Medicare outpatient hospitals by provider and service: methodology. U.S. Department of Health and Human Services. <https://data.cms.gov/provider-summary-by-type-of-service/medicare-outpatient-hospitals>. Reports services aggregated to Ambulatory Payment Classification (APC) groups, accessed via the CMS Data Catalog.
- Centers for Medicare & Medicaid Services (2024b). Medicare physician & other practitioners public use file: by provider and service. U.S. Department of Health and Human Services. <https://data.cms.gov/provider-summary-by-type-of-service/medicare-physician-other-practitioners>. Annual provider-by-HCPCS utilization file, accessed via the CMS Data Catalog.
- Clemens, J. and Gottlieb, J. D. (2014). Do physicians' financial incentives affect medical treatment and patient health? *American Economic Review*, 104(4):1320–1349.
- Dai, T. and Singh, S. (2025). Artificial intelligence on call: The physician's decision of whether to use ai in clinical practice. *Journal of Marketing Research*, 62(5):854–875.
- David, P. A. (1969). A contribution to the theory of diffusion. Memorandum 71, Stanford Center for Research in Economic Growth.
- Devaraj, S., Ow, T. T., and Kohli, R. (2013). Examining the impact of information technology and patient flow on healthcare performance: A theory of swift and even flow (TSEF) perspective. *Journal of Operations Management*, 31(4):181–192.
- Ding, P. and Li, F. (2019). A bracketing relationship between difference-in-differences and lagged-dependent-variable adjustment. *Political Analysis*, 27(4):605–615.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press, Princeton, NJ.
- Dobrzykowski, D. D. and Tarafdar, M. (2015). Understanding information exchange in healthcare operations: Evidence from hospitals and their supply chains. *Journal of Operations Management*, 36:201–214.
- Douglas, P. S., Hoffmann, U., Patel, M. R., Mark, D. B., Al-Khalidi, H. R., Cavanaugh, B., Cole, J., Dolor, R. J., Fordyce, C. B., Huang, M., Khan, M. A., Kosinski, A. S., Krucoff, M. W., Malhotra, V., Picard, M. H., Udelson, J. E., Velazquez, E. J., Yow, E., Cooper, L. S., and Lee, K. L. (2015). Outcomes of anatomical versus functional testing for coronary artery disease. *New England Journal of Medicine*, 372(14):1291–1300.
- Fairbairn, T. A., Nieman, K., Akasaka, T., Nørgaard, B. L., Berman, D. S., Raff, G., Hurwitz-Koweeek, L. M., Pontone, G., Kawasaki, T., Sand, N. P., Jensen, J. M., Amano, T., Poon, M., Øvrehus, K., Sonck, J., Rabbat, M., Mullen, S., De Bruyne, B., Rogers, C., Matsuo, H., Bax, J. J., Leipsic, J., and Patel, M. R. (2018). Real-world clinical utility and impact on clinical decision-making of coronary computed tomography angiography-derived fractional flow reserve: lessons from the ADVANCE Registry. *European Heart Journal*, 39(41):3701–3711.
- Faulder, T. I., Prematunga, K., Moloi, S. B., Faulder, L. E., Jones, R., and Moxon, J. V. (2024). Agreement of fractional flow reserve estimated by computed tomography with invasively measured fractional flow reserve: A systematic review and meta-analysis. *Journal of the American Heart Association*, 13(10).
- Galbraith, J. R. (1974). Organization design: An information processing view. *Interfaces*, 4(3):28–36.
- Gardner, J. W., Boyer, K. K., and Gray, J. V. (2014). Operational and strategic information processing: Complementing healthcare IT infrastructure. *Journal of Operations Management*, 33–34(1):123–139.
- Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.
- Goldratt, E. M. and Cox, J. (1984). *The Goal: A Process of Ongoing Improvement*. North River Press, Great Barrington, MA.
- Gruber, J. and Owings, M. (1996). Physician financial incentives and cesarean section delivery. *The RAND Journal of Economics*, 27(1):99.
- Gulati, M., Levy, P. D., Mukherjee, D., Amsterdam, E., Bhatt, D. L., Birtcher, K. K., Blankstein, R., Boyd, J., Bullock-Palmer, R. P., Conejo, T., Diercks, D. B., Gentile, F., Greenwood, J. P., Hess, E. P., Hollenberg, S. M., Jaber, W. A., Jneid, H., Joglar, J. A., Morrow, D. A., O'Connor, R. E., Ross, M. A., and Shaw, L. J. (2021). 2021 AHA/ACC/ASE/CHEST/SAEM/SCCT/SCMR guideline for the evaluation and diagnosis of chest pain. *Circulation*, 144(22).
- Heckman, J. J., Lalonde, R. J., and Smith, J. A. (1999). *The Economics and Econometrics of Active Labor Market Programs*, pages 1865–2097. Elsevier.
- Jing, A. B., Garg, N., Zhang, J., and Brown, J. J. (2025). AI solutions to the radiology workforce shortage. *npj Health Systems*, 2(1).
- Karshenas, M. and Stoneman, P. L. (1993). Rank, stock,

- order, and epidemic effects in the diffusion of new process technologies: An empirical model. *The RAND Journal of Economics*, 24(4):503.
- Koo, B.-K., Erglis, A., Doh, J.-H., et al. (2011). Diagnosis of ischemia-causing coronary stenoses by noninvasive fractional flow reserve computed from coronary computed tomographic angiograms. *Journal of the American College of Cardiology*, 58(19):1989–1997.
- Muehlematter, U. J., Bluethgen, C., and Vokinger, K. N. (2023). FDA-cleared artificial intelligence and machine learning-based medical devices and their 510(k) predicate networks. *The Lancet Digital Health*, 5(9):e618–e626.
- National Institute for Health and Care Excellence (2017). HeartFlow FFRCT for estimating fractional flow reserve from coronary CT angiography. Medical technologies guidance MTG32.
- Nørgaard, B. L., Leipsic, J., Gaur, S., et al. (2014). Diagnostic performance of noninvasive fractional flow reserve derived from coronary computed tomography angiography in suspected coronary artery disease. *Journal of the American College of Cardiology*, 63(12):1145–1155.
- Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591.
- Roemer, M. I. (1961). Bed supply and hospital utilization: a natural experiment. *Hospitals, J.A.H.A.*, 35:36–42.
- Roth, J., Sant’Anna, P. H., Bilinski, A., and Poe, J. (2023). What’s trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.
- Sivakumar, R., Lue, B., and Kundu, S. (2025). FDA approval of artificial intelligence and machine learning devices in radiology: A systematic review. *JAMA Network Open*, 8(11):e2542338.
- Srinivasan, R. and Swink, M. (2018). An investigation of visibility and flexibility as complements to supply chain analytics: An organizational information processing theory perspective. *Production and Operations Management*, 27(10):1849–1867.
- Sterman, J., Oliva, R., Linderman, K., and Bendoly, E. (2015). System dynamics perspectives and modeling opportunities for research in operations management. *Journal of Operations Management*, 39–40(1):1–5.
- Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.
- The SCOT-HEART Investigators (2018). Coronary CT angiography and 5-year risk of myocardial infarction. *New England Journal of Medicine*, 379(10):924–933.
- Thirumalai, S. and Devaraj, S. (2023). Mitigating the curse of complexity: The role of focus and the implications for costs of care. *Journal of Operations Management*, 70(1):157–179.
- Thirumalai, S., Devaraj, S., and Browning, T. R. (2024). Uncertainty in healthcare operations: How hospitals weather the perfect storm. *Journal of Operations Management*, 70(8):1194–1212.
- Tushman, M. L. and Nadler, D. A. (1978). Information processing as an integrating concept in organizational design. *Academy of Management Review*, 3(3):613–624.
- Zink, A., Chernew, M. E., and Neprash, H. T. (2026). Practice pattern changes after adoption of a diagnostic AI tool used in conjunction with cardiac imaging. *Health Affairs*, 45(3):238–245.