

The Radiology Distribution Test: Regulation and Model Drift in AI Valuation

A Valuation and Transaction Framework for Installed-Base Durability, Regulatory Scope, Distribution Economics and Continuous Model Validation

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Radiology artificial-intelligence companies can accumulate regulatory clearances, distributor agreements, hospital logos and installed software while producing materially different levels of clinical use and economic value. A product may be authorised for one intended use, deployed on a narrow modality or device, enabled at a site but rarely used, or embedded in a workflow without changing clinician action. Performance can also move after deployment as scanners, acquisition protocols, patient populations, prevalence, referral patterns, software versions and clinical practice change.

This paper develops a valuation and transaction framework for testing whether radiology AI distribution has become durable clinical infrastructure. It connects intended use, regulatory scope, analytical and clinical evidence, site activation, workflow integration, clinician action, reimbursement, customer cohorts, channel economics, post-market surveillance, model change, cybersecurity, quality systems and collected contribution. The method separates regulatory availability from commercial distribution, technical installation from productive use, and historical validation from continuing performance.

The analysis draws on current material from the US Food and Drug Administration, World Health Organization, International Medical Device Regulators Forum, National Institute for Health and Care Excellence, European Commission, Medicines and Healthcare products Regulatory Agency, Health Canada, standards organisations and clinical reporting initiatives.[1][2][3][4][5][6][7][8][9][10] These sources emphasise intended use, lifecycle risk management, representative evidence, clinical evaluation, transparency, predetermined change control, post-market monitoring and governance appropriate to the consequence of error.

A wholly hypothetical acquisition case illustrates a radiology AI company with USD 18.0 million of annual contracted revenue, USD 13.5 million of collected revenue, 28 contracted provider sites and 19 productively active sites. Every patient number, performance measure, activation rate, price, cost, probability, multiple and valuation amount is a management assumption created solely to demonstrate the method. None is a forecast, market benchmark, clinical claim or valuation opinion.

The paper concludes that an installed base supports durable value when the authorised commercial version works across the stated modality, device, population and workflow; clinicians act on the output; sites remain productively active; representative monitoring detects drift; change control preserves regulatory and clinical continuity; and distribution converts into retained contribution after channel, implementation and surveillance cost. Six figures and seven tables translate that conclusion into an evidence ladder, deployment map, cohort bridge, drift-control architecture, valuation framework, transaction protections and a 180-day programme.

Medical-device, clinical-practice, reimbursement, privacy, artificial-intelligence, competition, tax and corporate-law requirements vary by product and jurisdiction. Qualified clinical, regulatory, legal,

reimbursement, accounting and valuation specialists should determine the rules and evidence applicable to a specific technology and transaction.

JEL Classification: G24, G34, I11, I18, O33

Keywords: radiology artificial intelligence, diagnostic imaging, model drift, medical devices, distribution, valuation, mergers and acquisitions, post-market surveillance

JEL Classification: G24, G34, I11, I18, O33

Keywords: radiology artificial intelligence, diagnostic imaging, model drift, medical devices, distribution, valuation, mergers and acquisitions, post-market surveillance

1. Define the acquisition and valuation decision

The transaction question is whether the target has built a durable installed base or a collection of regulatory, commercial and technical positions that still require conversion. The buyer should identify which authorised products are contracted, installed, activated, used in eligible studies, acted upon by clinicians, renewed and profitable after deployment and surveillance cost.

The intended use defines the perimeter. The file should identify the clinical condition, target population, modality, acquisition device, image type, user, setting, model output, required action and consequence of error. A worklist-prioritisation tool differs from lesion detection, measurement, segmentation, diagnostic support or autonomous interpretation. Each use has distinct evidence, workflow, liability and economic requirements.

Distribution should be decomposed into rights, access and operating proof. A distributor agreement can create market access without customer demand. A hospital contract can exist before security approval, integration or go-live. An installation can remain outside the routine worklist. A technically active model can produce outputs that clinicians ignore. The valuation model should connect these states without treating one as evidence of another.

Value should be separated into collected stand-alone contribution, protected installed-base continuity, funded conversion, buyer-specific channel synergy and unproven option value. Regulatory authorisation, publications, partnerships and customer logos support only the links they directly evidence.


2. Build a distribution-to-value evidence ladder

The ladder begins with authorised scope and lawful market access. It then tests channel rights, contracted sites, technical deployment, productive clinical use, customer retention, contribution and post-market durability. Each layer supports a different decision and can fail independently.

Technical verification asks whether software, integrations and image pipelines operate as specified. Analytical validation tests output against an appropriate reference. Clinical validation assesses the relationship with the intended condition. Prospective utility examines whether using the output changes prioritisation, reporting, treatment or another material clinical or operational decision.

Implementation evidence measures security approval, interface performance, uptime, eligible-study capture, alert delivery, clinician review, turnaround time, overrides, training, exceptions and fallback. Economic evidence connects the live workflow to invoice, collection and delivery cost. Post-market evidence tests whether performance and utility persist as data and operating conditions change.

Figure 1. Proposed radiology-AI distribution-to-value ladder



Each layer supports a distinct decision; a higher commercial layer does not cure a missing clinical or technical control.

Table 1. Distribution evidence and valuation consequence

Evidence layer	Core question	Decision supported	Valuation consequence
authorised scope	which use, version and population?	lawful product perimeter	defines permitted opportunity
channel and contract	who can sell and who has bought?	commercial access	supports probability-weighted pipeline
technical deployment	does the product work in the site stack?	implementation readiness	supports go-live probability and cost
productive use	does the model affect routine work?	clinical adoption	supports active-site revenue
retention and contribution	do sites renew profitably?	commercial durability	supports central cash flow
post-market durability	does value survive change?	terminal economics	supports duration and multiple

Proposed framework; evidence requirements depend on intended use, risk, setting and jurisdiction.

3. Reconstruct the prospective study

The buyer should reconstruct the study from protocol to analysis rather than relying on the abstract or headline endpoint. It should identify registration, protocol version, statistical analysis plan, amendments, sites, enrolment, inclusion and exclusion criteria, comparator, sample-size assumptions, missing-data handling and stopping rules. Prospective registration and reporting frameworks such as SPIRIT-AI, CONSORT-AI and DECIDE-AI improve transparency but do not replace diligence.[32][33][34]

The model version is part of the tested intervention. The buyer should identify weights, thresholds, preprocessing, input devices, software dependencies and user interface. Any difference between the evaluated and commercial versions should be mapped to verification, validation and regulatory assessment. A product described by brand name can contain materially different clinical systems over time.

Site selection can create optimistic performance. Academic centres may have specialist staff, higher-quality data and implementation support unavailable to community sites. The team should compare trial

sites with the addressable market across prevalence, demographics, devices, workflows, staffing and care pathways. Excluded or failed cases belong in the economic population.

The comparator should represent actual care. A historical comparator can be affected by secular trends, coding, staffing or protocol changes. A silent deployment tests model performance without measuring behavioural response. A randomised workflow study can identify causal impact more strongly, but contamination, clinician learning and site effects still require analysis.

Table 2. Prospective-study reconstruction

Dimension	Evidence	Failure signal	Price response
protocol	registered plan and amendments	endpoint or population changed after data	reduce certainty and test sensitivity
model version	reproducible artefact and configuration	commercial version differs	exclude unsupported benefit
population	inclusion, exclusion and missing cases	narrow or selected cohort	cap addressable population
comparator	contemporaneous care pathway	weak or shifting baseline	lower causal confidence
sites	representative operating settings	expert-site dependence	price deployment cost and delay
analysis	prespecified methods and uncertainty	selective subgroup emphasis	use conservative central case

Proposed diligence catalogue; clinical and statistical specialists should determine materiality.

4. Read endpoints through the care pathway


Accuracy, sensitivity, specificity, area under a curve and calibration describe model behaviour under a defined population. Their economic meaning depends on prevalence, threshold, workflow and consequence. A small improvement in discrimination can have limited value when the operating point creates many false positives or when clinicians cannot act on the result.

The buyer should convert endpoints into a patient and work queue. True positives, false positives, false negatives and true negatives should be shown for the intended population. Each category can create investigation, treatment, delay, anxiety, harm, capacity use and cost. Confidence intervals and subgroup results should remain visible.

Process endpoints can be valuable when they connect to capacity or outcomes. Time to report, time to review, length of stay, cancelled procedure, avoided appointment and documentation time can support a purchasing decision. The bridge requires evidence that saved time is releasable, redeployed or converted into additional care.

Patient outcomes should match the mechanism and observation period. Mortality, morbidity, quality of life, diagnostic delay and adverse events differ in importance and statistical power. Composite endpoints should be disaggregated. A favourable surrogate can support value only where its relationship to meaningful outcomes is credible for the intended decision.

Figure 2. Proposed endpoint-to-economics chain



The chain requires evidence at every transition from output to collected economic value.

5. Reconstruct the installed base by version, modality and workflow


The installed-base ledger should identify customer, site, department, modality, scanner or acquisition device, image standard, integration, authorised intended use, model version, activation date, eligible studies, processed studies, delivered outputs, clinician interactions, incidents, contract, revenue, cost and current status. Aggregate site counts can conceal inactive, partial or unsupported deployments.

Representativeness should match the expansion plan. Performance can vary across computed tomography, magnetic resonance imaging, radiography, ultrasound, mammography and other modalities; within each modality it can vary by manufacturer, protocol, slice thickness, reconstruction, contrast, field strength and patient mix. A validation set drawn from one environment should not be assumed to cover another.

The ledger should retain retired and failed deployments. Exclusions reveal the cost and limits of distribution. A site may stall because of security review, interface complexity, inadequate study volume, workflow resistance, reimbursement, weak local ownership or model performance. These causes have different remedies and valuation effects.

Installed-base value arises from repeatable conversion. The buyer should measure elapsed time and cost from contract to security approval, interface, validation, go-live, productive use and renewal. A growing estate that requires bespoke engineering and clinical support at every site can produce revenue while weakening contribution and scale.

Figure 3. Hypothetical radiology-AI site conversion funnel



Counts are wholly hypothetical and demonstrate the distinction among contracted, deployed and productive sites.

6. Separate model performance from clinical utility


A model can be accurate and clinically irrelevant. It may identify information already available to the clinician, arrive after a decision, recommend an unavailable action or add alerts to a saturated workflow. Clinical utility requires a changed decision or action that is appropriate for the patient and feasible for the service.

The buyer should measure receipt, opening, comprehension, agreement, override, action and timing. Each step forms an adoption funnel. An apparent low adoption rate can signal distrust, poor placement, inadequate training, weak relevance or appropriate rejection. Override reasons and subsequent outcomes should be captured rather than classified automatically as user error.

Human factors deserve the same diligence as the model. Interface design, alert wording, prioritisation, explanation, training, escalation and downtime procedures can alter safety and performance. The FDA and WHO emphasise transparency, user needs and human oversight across the lifecycle.[1][3] The tested system includes people, process and technology.

Capacity can reverse value. A detection tool that finds more patients may increase imaging, biopsy, specialist review, treatment or follow-up demand. If downstream capacity is unavailable, the product can create longer queues or displace higher-value care. The economic model should include the complete pathway and any bottleneck investment.

Figure 4. Proposed radiology-AI productive-use funnel



Hypothetical funnel; percentages are illustrative management assumptions and are not clinical benchmarks.

7. Map regulatory scope to every distributed configuration

Regulatory status should be mapped at product, intended-use, model-version and jurisdiction level. The approved or cleared description, indications, contraindications, users, inputs, outputs and limitations should be reconciled with contracts, demonstrations, training, user interfaces and actual workflow. Commercial language that extends beyond authorised scope creates liability and valuation risk.

Distribution adds configuration risk. The deployed model may interact with different acquisition devices, picture archiving and communication systems, radiology information systems, viewers, worklists and reporting tools. The buyer should identify which configurations were included in verification and validation, which are supported through documented equivalence, and which require further evidence or regulatory assessment.

Model updates, threshold changes, retraining, data-pipeline changes and third-party components should be traced through quality and change control. The FDA's lifecycle and predetermined-change-control guidance, EU medical-device framework, UK programme and international good machine-learning-practice principles provide relevant control boundaries.[1][2][5][6][7][8]

Regulatory continuity belongs in the transaction model. A change in ownership, legal manufacturer, authorised representative, quality system, hosting, supplier or product architecture can create filing, notification, certification and customer obligations. The closing and integration plan should preserve the evidence and roles required to keep the commercial version lawful and supported.

8. Price reimbursement and budget ownership

Regulatory authorisation permits a product to be marketed under defined conditions; it does not create reimbursement or a purchasing budget. The buyer should identify who pays, the covered service, coding, price, evidence requirement, utilisation control, contracting route and renewal decision. Payment can arise through a dedicated code, bundled payment, hospital budget, population contract, subscription or shared-savings arrangement.

The economic beneficiary can differ from the buyer. A hospital department may bear software and implementation cost while savings accrue to a payer, another department or future period. The transaction

model should show the value flow and the contracting mechanism that returns an appropriate share to the vendor.

NICE's evidence standards framework connects clinical effectiveness, real-world performance, implementation, budget impact and value to commissioning decisions.[4][26] Early-use HealthTech guidance can support managed adoption with evidence generation, but conditional access should be distinguished from durable routine purchasing.[27][28]

US reimbursement can vary across setting, payer, code and service. A code does not guarantee coverage, payment amount, utilisation or vendor capture. The buyer should reconcile claims, denials, appeals, contract allowances and cash by site and payer. Provider-paid models require the same discipline around budget ownership and renewal.

Table 3. Reimbursement and budget map

Route	Buyer	Evidence needed	Main valuation risk
dedicated reimbursement	provider or supplier	coverage, coding and claim acceptance	code without durable coverage
bundled payment	provider system	pathway saving and capacity	benefit absorbed by bundle
departmental budget	clinical or operating unit	local utility and budget impact	annual discretionary renewal
enterprise licence	health system	multi-site performance and governance	slow procurement and deployment
population contract	payer or risk-bearing provider	outcome and total-cost evidence	attribution and contract duration
evidence-generation access	public or provider programme	protocol and reporting	temporary adoption mistaken for scale

Proposed map; payment rules and contracting routes require jurisdiction-specific verification.

9. Rebuild contracted, recognised and collected revenue

Clinical AI contracts can mix software licences, per-study fees, implementation, integration, support, research services and outcome-linked payments. The buyer should reconcile signed contract value, deployed sites, eligible volume, recognised revenue, invoicing, deductions and cash. A multi-year headline contract can contain termination rights, pilots, minimums that are not enforced or future modules that remain unapproved.

Contracted annual recurring revenue should be segmented by live, implementing and unstarted sites. A live site should have the intended model version in production, trained users, functioning integration, monitored activity and an accepted invoice route. Implementation revenue should remain separate from recurring licence or usage economics.

The hypothetical case reports USD 18.0 million of annual contracted revenue and USD 13.5 million of collected revenue across 28 provider sites. The gap comprises deployment timing, minimum-volume shortfall, invoice dispute and receivable delay. Every amount is hypothetical and created to demonstrate reconciliation.

Collections should be traced to bank receipt. Healthcare procurement, purchase orders, acceptance, coding, claim adjudication and budget cycles can create material delay. Bad debt, credit notes, rebates, free extensions and implementation concessions should reduce the appropriate cohort economics.

10. Build site and study cohorts

Site cohorts should begin at contract and progress through security approval, technical installation, validation, activation, productive use, renewal and collection. Study cohorts should begin with eligible images and trace ingestion, model processing, output delivery, clinician review, action, report, downstream care and any measured outcome.

A contracted site is not an active site. An installed site can process little eligible volume. A processing site can show weak clinician interaction. A renewal can be supported by a broader enterprise agreement rather than product utility. The buyer should retain each evidence state and reconcile movements by period.

Productive use needs a defined threshold appropriate to the intended use. Measures can include eligible-study capture, successful processing, output latency, worklist availability, clinician opening, acknowledgement, action, override, report incorporation and time to decision. The threshold should be stable enough to compare cohorts and transparent enough for customers and investment committees.

Retention should be examined by site, modality, workflow, channel, contract type and model version. A distributor can keep a contract nominally alive while end-user activity declines. Direct enterprise revenue and channel revenue can have different visibility, support cost, renewal rights and cash timing.

11. Reconstruct unit economics after implementation


Revenue should be reduced by variable and directly attributable cost. Clinical data access, cloud, inference, integration interfaces, cybersecurity, annotation, quality assurance, clinical support, regulatory maintenance, post-market surveillance and customer success can be material. A software gross margin that excludes these functions can overstate contribution.

Implementation cost should include interface configuration, data mapping, validation, workflow design, information governance, training, parallel operation and remediation. Some cost is reusable; some repeats at every site. The buyer should separate platform, country, health-system and site-specific work.

The hypothetical case collects USD 13.5 million. Data, cloud and model operations cost USD 1.7 million; site integration and support cost USD 2.4 million; clinical quality, regulatory and monitoring cost USD 1.5 million; sales, procurement and partner cost USD 1.2 million; and incident, insurance and security cost USD 0.7 million. Contribution before central cost, tax and capital is USD 6.0 million.

Each amount is hypothetical. The example does not claim a representative margin. It demonstrates that prospective evidence creates value only when deployment, governance and monitoring cost remain inside the commercial model.

Figure 5. Hypothetical radiology-AI contribution bridge



Wholly hypothetical USD millions; central cost, tax and capital remain outside the displayed contribution.

Table 4. Clinical-AI cohort economics

Measure	Numerator	Denominator	Diligence use
input usability	cases processed without invalid input	eligible cases	data and integration quality
output review	outputs opened by intended user	delivered outputs	workflow adoption
appropriate action	supported actions completed	reviewed outputs	clinical utility
active-site retention	productive sites retained	productive sites at opening	operating durability
contribution per case	collected revenue less attributable cost	completed pathways	scalable economics
cash conversion	cash collected	recognised revenue	procurement and payer quality

Proposed structure; financial and clinical populations should reconcile to governed source records.

12. Quantify evidence-generation cost and dilution

Additional evidence is an investment. The valuation model should include protocol design, sites, data, clinical operations, statistics, monitoring, quality, regulatory interaction and publication. It should also include time to enrol, site activation, observation and analysis. The company may require financing before the evidence event.

The financing path can dilute shareholders or add preference. A buyer pricing the company before a pivotal study should model the capital needed to reach the readout and the rights attached to that capital. Enterprise value and equity proceeds should not be conflated.

Evidence cost should be linked to the decision it can unlock. A large randomised study can be wasteful when the immediate uncertainty is integration feasibility or payer willingness. A small pilot can be insufficient when the claim requires patient-outcome evidence. The next study should target the uncertainty with the greatest expected effect on value.

Study failure should be decomposed. Recruitment delay, missing data, poor model performance, low clinician adoption, inadequate capacity and weak economics require different remediation. Some preserve the core asset; others challenge the product thesis.

The board should maintain an evidence-financing schedule. Each work package should state the decision it supports, start and completion date, cash requirement, dependency, success criterion and residual uncertainty. This makes the capital path auditable and prevents a sequence of loosely defined studies from consuming cash without changing the commercial decision.

Partnership terms can transfer evidence cost while surrendering economics or control. A strategic hospital, manufacturer or payer may fund a study in exchange for exclusivity, data rights, distribution, preferred pricing or an acquisition option. The buyer should value the funding benefit alongside the restriction on addressable market and future bargaining power.

Publication timing can differ from the valuation event. A prespecified analysis may support a regulatory or purchasing decision before peer-reviewed publication, while an attractive publication can describe an obsolete version. The diligence room should contain the governed underlying evidence, analysis code, approvals and limitations rather than treating journal acceptance as the sole quality gate.

13. Design a health-economic model

The comparator should be current practice at the purchasing setting. It should include tests, clinician time, treatment, complications, follow-up, administration and capacity. The AI pathway should include software, implementation, training, monitoring, false-positive work-up, false-negative consequence and downstream care.

Time horizon should match outcomes and payment. A hospital may focus on annual budget impact while a payer values multi-year outcomes. Discounted cost per outcome, budget impact, capacity and cash should be presented separately. A favourable cost-effectiveness result does not prove affordability or vendor revenue.

The model should distinguish resource release from theoretical time saving. Five minutes saved per case creates cash value only if staffing, overtime, outsourcing, throughput or avoided delay changes. Capacity released for additional care can be valuable even without cost reduction, but it requires demand and operational execution.

Uncertainty should be explicit. Eligible volume, adoption, prevalence, diagnostic performance, action rate, outcome effect, pathway cost, price and persistence should be varied. Structural assumptions should receive scenarios rather than narrow confidence intervals.

Table 5. Health-economic evidence architecture

Layer	Observed evidence	Modelled transition	Value test
population	eligible and usable cases	addressable volume	avoid prevalence inflation
performance	threshold results and uncertainty	errors by case type	price harm and work-up
action	reviewed outputs and overrides	changed care	separate output from utility
outcome	patient and capacity measures	longer-term effect	match time horizon
resource	staff, test and pathway use	released or added capacity	distinguish time from cash
payment	contract, claim and collection	renewal and price	connect benefit to vendor cash

Proposed architecture; clinical and economic inputs require source-specific validation.

14. Value the installed base by evidence state

The valuation should begin with collected contribution from productively active and retained sites. Contracted sites that have not gone live belong in a conversion schedule with explicit probability, time, deployment cost and customer obligations. New territories, modalities and autonomous uses belong in option value until regulatory, clinical and commercial evidence matures.

The hypothetical case contains 28 contracted sites, 23 technically live sites, 19 productively active sites and 16 sites with at least twelve months of seasoned use. Contracted annual revenue is USD 18.0 million and collected revenue is USD 13.5 million. Every figure is illustrative. The bridge demonstrates why a company cannot value all contracts as equally mature recurring revenue.

Active-site contribution should deduct channel commission, interface maintenance, cloud and inference cost, clinical support, customer success, quality, cybersecurity, post-market surveillance and customer-specific development. A higher gross revenue multiple can be misleading when each new site needs bespoke work or mature sites retain extensive support.

Comparable-company and transaction evidence should be normalised for regulated scope, product mix, active-site penetration, growth, retention, margin, customer concentration and capital needs. Discounted cash flow should make conversion and drift assumptions visible. Replacement cost can support an asset analysis but does not prove that an installed base will renew or create cash.

15. Apply a model-drift discount transparently

Model drift is a change in the relationship among inputs, outputs, reference standards, clinical decisions and outcomes. Input drift can arise from devices, protocols, compression, workflow or population. Prevalence and referral change can alter predictive values. Label or practice change can alter the apparent reference. Model and software updates can change performance or interaction.

The buyer should quantify drift through a versioned monitoring plan. Measures can include data-quality failures, input-distribution changes, output-rate changes, subgroup performance, calibration, sensitivity, specificity, false-priority burden, overrides, clinician action, incidents and downstream outcomes. Thresholds should connect to investigation, restriction, rollback, revalidation and regulatory assessment.

The valuation discount should follow cash-flow causality. Drift can reduce eligible volume, productive-site retention, price, renewal, margin or expansion probability; it can also increase monitoring, remediation, evidence and regulatory cost. Once those effects are included in cash flow, an additional multiple discount should cover only residual uncertainty.

A strong monitoring system can reduce uncertainty. The system needs representative data, stable denominators, ground truth or adjudication where appropriate, site participation, version lineage, accountable review and authority to act. A dashboard without an operating response does not preserve value.

16. Test sensitivities and downside cases

Sensitivity should expose variables management can influence and variables controlled by the environment. Eligible volume, input usability, clinician review, action, price, deployment cost and renewal can be operationally addressed. Regulatory timing, payer coverage and market multiples may have less direct control.

Downside cases should include a failed new-site replication, reimbursement delay, material performance drift, integration outage, cybersecurity incident, regulatory restriction and loss of a clinical champion. The model should show liquidity needs and equity proceeds as well as enterprise value.

The hypothetical contribution sensitivity varies collected revenue and attributable operating cost. It excludes central cost, tax, capital and transaction financing. No cell is a forecast or benchmark.

Table 6. Hypothetical annual contribution sensitivity

Collected revenue; USDm	Attributable cost USD 6.5m	USD 7.5m	USD 8.5m	USD 9.5m
12.5	6.0	5.0	4.0	3.0
13.5	7.0	6.0	5.0	4.0
14.5	8.0	7.0	6.0	5.0
15.5	9.0	8.0	7.0	6.0

Wholly hypothetical USD millions; no cell is a forecast or market benchmark.

17. Translate evidence into transaction protections

Representations can address regulatory status, intended use, study conduct, data rights, model ownership, clinical claims, quality systems, cybersecurity, incidents, reimbursement, contracts and financial metrics. Definitions should match the evidence populations and deployed versions.

Conditions can require regulatory approval, change-of-control consent, transfer of essential data and model rights, delivery of reproducible validation, closure of a critical quality gap or execution of a material customer contract. Conditions should be reserved for matters that prevent lawful ownership or safe operation.

Interim covenants should protect the evidence base between signing and closing. Model updates, threshold changes, new claims, protocol amendments, material incidents, customer concessions and loss of clinical or regulatory personnel can alter value. Ordinary-course operation should continue within defined change control.

Escrow, indemnity and insurance should match enforceable exposure. Known evidence gaps and forward performance may receive limited warranty-insurance coverage. The buyer should model uninsured exposure, remediation cost and recovery timing.

Contingent consideration can align payment with effective authorisation, successful prospective replication, retained productive sites, reimbursement, collected contribution or a defined patient outcome. The measure needs a reproducible population, version, observation period, cost allocation, audit right and dispute route.

Milestones should reward evidence within the seller's influence while preserving safe operation. A payment tied only to regulatory submission may reward activity rather than acceptance or commercial value. A payment tied to revenue can encourage discounting, unsupported expansion or underinvestment in monitoring. Balanced milestones can combine permission, productive use, quality and collected contribution.

Operating covenants during an earn-out should define the buyer's obligations without freezing integration. The parties should agree minimum resources, commercial discretion, model change, customer allocation, pricing authority and access to evidence. An independent expert process can resolve technical or clinical disputes more effectively than a purely financial definition.

The transaction committee should map each material gap to one primary response. A weakness reflected in forecast volume should not automatically be deducted again from price and protected by a full-value earn-out. The map should show where risk is priced, retained, insured, remediated or transferred.

Table 7. Evidence-to-protection matrix

Evidence gap	Price response	Protection	Release evidence
narrow prospective cohort	probability-weight expansion	milestone consideration	independent replication
commercial version divergence	exclude unsupported benefit	closing validation	reproducible matched performance
pending reimbursement	defer payment-route value	earn-out or option	coverage and collected claims
weak site adoption	reduce volume and margin	retention mechanism	mature productive cohorts
quality-system gap	funded remediation deduction	condition, escrow and covenant	tested closure
model-drift risk	terminal-value adjustment	monitoring and change control	stable post-market evidence
key clinical dependency	continuity discount	retention and transfer plan	documented repeatable governance

Proposed matrix; legal drafting and remedies remain transaction-specific.

18. Design integration around clinical continuity

Integration can invalidate the evidence that supported value. Hosting, data mapping, interface, model version, threshold, user interface, workflow, support and quality responsibility may change. Each change should be mapped to risk assessment, verification, validation, regulatory review and customer acceptance.

The integration architecture should distinguish preserve, connect, migrate and retire. Preserve applies where the target's regulated platform and quality system remain necessary. Connect uses governed interfaces while operating systems remain separate. Migrate moves controlled cohorts after matched testing. Retire follows evidence that records, obligations and safe fallback have transferred.

Data migration should preserve provenance, consent, time, units, coding, missingness and deletion state. Normalisation can silently alter clinical meaning. Record-level reconciliation and clinically meaningful test cases should precede production use.

Model migration is a controlled clinical change. Old and new systems should process matched cases where feasible. Differences should be classified by input, preprocessing, model, threshold, interface and user response. The migration gate should require acceptable clinical, operational, security and financial results.

Customer continuity comes first. Clinicians need access, training, escalation and downtime procedures. Patients should not experience an unannounced reduction in care. Cost synergy should be released after replacement controls and operating capacity are proven.

19. Govern model change and post-market performance

The model register should connect each commercial version to training and validation data, intended use, regulatory record, software and supplier dependencies, supported configurations, deployed sites, monitoring, incidents, changes and retirement. The buyer should be able to reproduce which version produced an output at a specific site and time.

Change control should classify updates by clinical and regulatory consequence. A cybersecurity patch, interface change, threshold adjustment, model retraining, expanded population and new autonomous action require different evidence and approvals. Predetermined change-control plans can define bounded

modifications where accepted by the regulator; they do not remove the need for implementation, monitoring and documentation.

Post-market surveillance should combine technical, clinical and commercial evidence. Performance statistics without workflow context can miss low adoption. Usage without quality evidence can scale harm. Renewal without contribution can scale economic loss. The board should receive a connected view of version, use, performance, action, incidents, retention, cost and cash.

Customer communication and fallback belong in the control system. Sites should know which version is live, material limitations, relevant changes, monitoring responsibilities and escalation routes. The product should preserve a safe degraded mode, rollback or alternative workflow when the AI service is unavailable or restricted.

20. Retain clinical, regulatory and implementation capability

Clinical AI value often resides in people who understand intended use, study design, quality systems, clinical workflow, customers and regulators. The buyer should map critical roles, authority, relationships, documentation and succession. Employment title is a weak proxy for actual dependency.

Retention should focus on capability transfer. Transition plans should document model design, data decisions, thresholds, protocol history, regulatory interactions, incidents, customer configurations and reimbursement assumptions. Access and signing authority should move through controlled processes.

Clinical champions can support adoption while creating concentration risk. The buyer should identify whether evidence and customer relationships depend on a founder clinician, principal investigator or single site. Independent governance and repeatable deployment reduce that dependency.

The combined organisation needs named owners for each product, model, clinical claim, regulatory file, customer pathway and outcome. Centralisation can improve consistency while weakening local response if authority and workflow knowledge are removed too quickly.

21. Allocate distribution synergy with evidence

Distribution synergy can arise from access to additional hospitals, imaging networks, modalities, geographic channels, enterprise contracts or complementary workflows. The buyer should identify the precise customer, authorised use, integration, sales route, clinical sponsor, budget and time to productive deployment.

Cross-selling should not be valued from customer overlap alone. The target product may need new evidence, security review, interface work, procurement, training or reimbursement. A parent company's account relationship can accelerate access while clinical and regulatory conversion remains incomplete.

Platform synergy can reduce hosting, security, integration or support cost where systems can be combined without disrupting authorised configurations or customer obligations. Data synergy requires rights, compatibility, relevance and governance. Broader data can improve monitoring and can also reveal heterogeneity that reduces headline performance.

The synergy ledger should record gross opportunity, enabling investment, probability, timing, owner, customer evidence, regulatory dependency, clinical risk and collected result. Seller stand-alone value and buyer-specific synergy should remain separate.

22. Execute an evidence-gated 180-day programme


Days one to thirty should preserve authorised versions, quality records, integrations, monitoring, customer service, incident response and key clinical, regulatory and engineering personnel. The buyer should freeze uncontrolled change and establish decision rights.

Days thirty to seventy should reconcile the installed-base ledger, reproduce material validation results, map deployed configurations, trace channel and customer rights, normalise site cohorts and verify revenue, collection and delivery cost. Unsupported configurations and high-risk drift signals should enter controlled review.

Days seventy to one hundred and twenty should remediate priority quality, data, security, interface and monitoring gaps. The combined firm should pilot integrations on reversible site cohorts with predefined clinical, operational and economic acceptance criteria.

Days one hundred and twenty to one hundred and eighty should season productive-use, performance, renewal and contribution evidence; complete approved migrations; test fallback and incident response; and release contingent value only after agreed gates pass.

Figure 6. Proposed evidence-gated 180-day radiology-AI programme



Timing should follow clinical, regulatory, customer, channel and technology constraints.

23. Decision and conclusion

Radiology AI distribution creates durable enterprise value when market access converts into a governed installed base that remains clinically useful and economically retained. The chain begins with authorised intended use and representative evidence. It continues through supported configurations, technical activation, productive workflow use, clinician action, customer renewal, contribution and collected cash.

The installed-base ledger prevents commercial labels from replacing operating evidence. Regulatory clearances, distributor agreements, contracts, installations and active use occupy different states. The buyer should value each state according to conversion probability, time, cost, rights and risk.

Model drift affects both clinical and financial duration. Representative monitoring should detect relevant changes in inputs, performance, clinician interaction and outcomes; change control should connect the signal to investigation, restriction, rollback, revalidation and regulatory assessment. A controlled lifecycle supports retention and terminal value.

The acquisition model should price current collected contribution centrally, fund verified conversion, probability-weight expansion and preserve buyer-specific synergy outside seller stand-alone value.

Consideration, protections and the ownership plan should follow identified evidence gaps and controllable risks.

The final test is whether a buyer can reproduce how each distributed configuration turns an eligible image into a governed output, an appropriate clinical action, a retained customer and collected contribution while remaining within its authorised scope. Evidence that survives this test supports value. A break identifies remediation, structure, price adjustment or a reason to stop.

Frequently asked questions

Q: Does regulatory clearance establish the value of a radiology AI installed base? A: It establishes a defined regulatory position. Enterprise value also depends on supported configurations, productive site use, performance durability, customer retention, contribution and collected cash.

Q: What is the difference between installation and productive clinical use? A: Installation means that the product is technically available. Productive use means that eligible studies are processed, outputs reach the intended workflow, clinicians use them appropriately and the site sustains the process.

Q: Is regulatory authorisation the same as reimbursement? A: No. Authorisation addresses lawful marketing for a defined use. Reimbursement or budget ownership requires a separate payer, coding, coverage, procurement or provider decision and evidence of value.

Q: How should a buyer value a large number of hospital logos? A: The buyer should reconcile logos to contracts, installations, eligible-study volume, productive use, renewal, contribution and cash. Each state supports a different part of value.

Q: How should model drift affect valuation? A: The buyer should model the effect of drift on productive volume, retention, price, margin, monitoring, remediation and regulatory cost. Residual uncertainty can affect duration or transaction structure.

Q: Which metric matters most after deployment? A: No single metric is sufficient. A governed chain should connect eligible cases, usable inputs, outputs, clinician action, patient or capacity outcomes, invoices, attributable cost and collected cash.

Q: Can time saved by clinical AI be treated as cash saving? A: Only when staffing, overtime, outsourcing, throughput or another measurable resource changes. Theoretical minutes saved can still create capacity value, but demand and operational execution must be demonstrated.

Q: What should happen during the first 180 days after a radiology AI acquisition? A: Preserve clinical continuity, reproduce evidence, map versions and rights, remediate material controls, pilot reversible integration, season outcomes and release value only after agreed clinical and commercial gates pass.

References

- [1] US Food and Drug Administration, Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing>
- [2] US Food and Drug Administration, Predetermined Change Control Plan for AI-Enabled Device Software Functions <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial>
- [3] World Health Organization, Regulatory considerations on artificial intelligence for health <https://www.who.int/publications/i/item/9789240078871>
- [4] National Institute for Health and Care Excellence, Evidence standards framework for digital health technologies <https://www.nice.org.uk/what-nice-does/digital-health/evidence-standards-framework-esf-for-digital-health-technologies>
- [5] European Commission, MDCG endorsed guidance for medical devices https://health.ec.europa.eu/medical-devices-sector/new-regulations/guidance-mdcg-endorsed-documents-and-other-guidance_en
- [6] European Commission, MDCG 2025-6 interplay between medical-device regulation and the AI Act https://health.ec.europa.eu/system/files/2025-06/mdcg_2025-6_en.pdf
- [7] Medicines and Healthcare products Regulatory Agency, Software and AI as a Medical Device Change Programme <https://www.gov.uk/government/publications/software-and-ai-as-a-medical-device-change-programme>

- [8] US Food and Drug Administration, Health Canada and MHRA, Good Machine Learning Practice for Medical Device Development <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles>
- [9] International Medical Device Regulators Forum, Software as a Medical Device Clinical Evaluation <https://www.imdrf.org/documents/software-medical-device-samd-clinical-evaluation>
- [10] International Medical Device Regulators Forum, Machine Learning-enabled Medical Devices Key Terms and Definitions <https://www.imdrf.org/documents/machine-learning-enabled-medical-devices-key-terms-and-definitions>
- [11] US Food and Drug Administration, Clinical Decision Support Software <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>
- [12] US Food and Drug Administration, Cybersecurity in Medical Devices <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions>
- [13] US Food and Drug Administration, AI and Machine Learning Enabled Medical Devices <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices>
- [14] US Food and Drug Administration, Performance Evaluation Methods for Evolving AI-Enabled Medical Devices <https://www.fda.gov/medical-devices/medical-device-regulatory-science-research-programs-conducted-osel/performance-evaluation-methods-evolving-artificial-intelligence-ai-enabled-medical-devices>
- [15] World Health Organization, Ethics and governance of artificial intelligence for health <https://www.who.int/publications/i/item/9789240029200>
- [16] World Health Organization, Ethics and governance of AI for health: large multi-modal models <https://www.who.int/publications/i/item/9789240084759>
- [17] World Health Organization, Guidance on large multi-modal models for health <https://www.who.int/news/item/18-01-2024-who-releases-ai-ethics-and-governance-guidance-for-large-multi-modal-models>
- [18] World Health Organization, Artificial intelligence for health <https://www.who.int/publications/m/item/artificial-intelligence-for-health>
- [19] European Commission, Medical Device Regulation <https://eur-lex.europa.eu/eli/reg/2017/745/oj>
- [20] European Commission, In Vitro Diagnostic Medical Devices Regulation <https://eur-lex.europa.eu/eli/reg/2017/746/oj>
- [21] European Union, Artificial Intelligence Act <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [22] European Union, General Data Protection Regulation <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [23] European Commission, MDCG 2020-1 guidance on clinical evaluation of medical device software https://health.ec.europa.eu/system/files/2020-09/md_mdcg_2020_1_guidance_clinic_eva_md_software_en_0.pdf
- [24] European Commission, MDCG 2019-11 qualification and classification of software https://health.ec.europa.eu/system/files/2020-09/md_mdcg_2019_11_guidance_qualification_classification_software_en_0.pdf
- [25] European Medicines Agency, artificial intelligence in the medicinal product lifecycle <https://www.ema.europa.eu/en/use-artificial-intelligence-ai-medicinal-product-lifecycle-scientific-guideline>
- [26] National Institute for Health and Care Excellence, How to meet the digital-health evidence standards <https://www.nice.org.uk/corporate/ecd7/chapter/how-to-meet-the-standards>
- [27] National Institute for Health and Care Excellence, Early-use HealthTech guidance assessments <https://www.nice.org.uk/process/pmg48/chapter/early-use-healthtech-guidance-assessments>
- [28] National Institute for Health and Care Excellence, Evidence generation for AI fracture detection <https://www.nice.org.uk/guidance/htg739/resources/evidence-generation-plan-for-artificial-intelligence-ai-technologies-to-help-detect-fractures-on-xrays-in-urgent-care-13618426285/chapter/approach-to-evidence-generation>
- [29] National Institute for Health and Care Excellence, HealthTech programme methods <https://www.nice.org.uk/process/pmg48/chapter/methods-for-guidance-produced-in-the-nice-healthtech-programme>
- [30] National Institute for Health and Care Excellence, Budget impact test for digital health technologies <https://www.nice.org.uk/corporate/ecd7/chapter/how-to-meet-the-standards>
- [31] Health Canada, Pre-market guidance for machine learning-enabled medical devices <https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/application-information/guidance-documents/pre-market-machine-learning-enabled-medical-devices.html>
- [32] SPIRIT-AI extension, protocol guidance for clinical trials involving artificial intelligence <https://www.nature.com/articles/s41591-020-1037-7>
- [33] CONSORT-AI extension, reporting clinical trials involving artificial intelligence <https://www.nature.com/articles/s41591-020-1034-x>
- [34] DECIDE-AI, reporting early-stage clinical evaluation of AI decision support <https://www.nature.com/articles/s41591-022-01772-9>
- [35] TRIPOD+AI statement, reporting prediction models using regression or machine learning <https://www.bmj.com/content/385/bmj-2023-078378>
- [36] PROBAST+AI, risk of bias assessment for prediction models <https://www.probast.org/>
- [37] EQUATOR Network, artificial intelligence reporting guidelines <https://www.equator-network.org/>
- [38] National Institute of Standards and Technology, AI Risk Management Framework <https://www.nist.gov/itl/ai-risk-management-framework>
- [39] National Institute of Standards and Technology, Cybersecurity Framework 2.0 <https://www.nist.gov/cyberframework>
- [40] International Organization for Standardization, ISO 13485 medical-device quality management <https://www.iso.org/standard/59752.html>
- [41] International Organization for Standardization, ISO 14971 medical-device risk management <https://www.iso.org/standard/72704.html>
- [42] International Organization for Standardization, ISO IEC 42001 AI management systems <https://www.iso.org/standard/81230.html>
- [43] International Organization for Standardization, IEC 62304 medical-device software lifecycle <https://www.iso.org/standard/64686.html>
- [44] International Organization for Standardization, IEC 62366-1 usability engineering <https://www.iso.org/standard/63179.html>
- [45] US Centers for Medicare and Medicaid Services, Medicare Coverage Database <https://www.cms.gov/medicare-coverage-database/>
- [46] US Centers for Medicare and Medicaid Services, coverage criteria and use of algorithms in Medicare Advantage <https://www.cms.gov/files/document/hpms-memo-faq-coverage-criteria-and-utilization-management-cms-4201-f-02-6-2024-pdf.pdf>
- [47] Organisation for Economic Co-operation and Development, AI principles <https://oecd.ai/en/ai-principles>
- [48] International Valuation Standards Council, International Valuation Standards <https://ivsc.org/standards/>
- [49] IFRS Foundation, IFRS 3 Business Combinations <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [50] IFRS Foundation, IAS 38 Intangible Assets <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.