BOOK EIGHT · THE AI ECONOMY MONETIZATION SERIES

The Open-Claw Effect

The widening gap between what AI can do and what companies are able to charge for it

The open claw opens silently. Every quarter you wait, it widens. The question is not whether you have an open-claw problem. It is how big yours already is.

Framework F22: The Claw Dynamics Model + The Four Closers

PREFACE

Picture a Lobster Claw

Why the gap between AI capability and revenue capture is the defining commercial challenge of the AI era.

Picture a lobster claw.

One arm represents AI capability — what the technology can actually do. It is moving upward rapidly. Each quarter brings new models with new reasoning powers, new speed, new cost efficiency, new capability frontiers that the previous quarter's models could not approach. The trajectory is exponential.

The other arm represents revenue capture — what companies are actually charging for those capabilities. It is moving upward too — but more slowly. Contracts signed in 2022 are still running at 2022 pricing. Customers using AI-powered features that are dramatically better than they were at contract signing are paying 2022 prices for 2025 capabilities. Sales teams are still selling the products they understand, with the pricing models that were negotiated two years ago, to procurement processes that move at human speed.

The gap between the two arms — the open space between what the AI can do and what the business charges for it — is the open claw. It widens every quarter. And it is widening right now, at most AI companies, in ways that their boards and investors have not fully quantified.

This book is about that gap: what causes it, how to measure it, and — most importantly — how to close it.

The open-claw effect is not a flaw in the AI economy. It is an inevitable consequence of the speed mismatch between technology advancement and commercial adaptation. AI capability moves at the pace of a competitive technology race. Commercial models move at the pace of enterprise procurement, contract cycles, and organizational change management. The gap is structural.

But structural does not mean inevitable. Companies that understand the dynamics of the open claw — that diagnose their gap, build the instruments to measure it, and implement the commercial strategies to close it — will capture a disproportionate share of the value their AI creates. Companies that allow the claw to continue opening will watch the return on their AI investment steadily compress as the gap between capability and capture widens.

The companies in the first category are building what will be the most valuable businesses of the next decade. The companies in the second category are funding their competitors' advantage with every quarter of delay.

This book gives you the framework, the measurement methodology, and the four specific strategies to move from the second category to the first.

PART ONE

The Claw Opens

The discovery. The four forces. Where the gap is widest. The claw as enabler.

CHAPTER ONE

The Discovery: Your AI Is Worth More Than You Are Charging For It

How to diagnose your open-claw gap. The measurement framework. What a large gap means strategically.

There is a specific moment when an AI company's leadership first truly confronts its open-claw problem. It usually arrives not as a strategic insight but as a jarring juxtaposition of two numbers: what the AI is demonstrably capable of doing, and what the company is actually charging for it.

Consider what happened at a mid-market legal AI company in late 2023. Their AI had, over the course of two years of continuous improvement, evolved from a contract flagging tool — it identified potentially problematic clauses and flagged them for attorney review — to a contract intelligence platform that could identify issues, suggest alternative language, assess risk across multiple dimensions, benchmark terms against market standards, and generate comprehensive review memoranda indistinguishable in quality from those produced by junior associates at major law firms.

The AI had improved dramatically. The price had not. The company was still charging $150 per contract review — the price established when the product was a clause-flagging tool. The current product, by any reasonable assessment of the value it delivered, was worth $600–800 per review in the context of what law firms and corporate legal departments were paying for equivalent junior associate work.

The gap — $150 charged versus $600–800 value delivered — was the open claw in its most visible form. The company was delivering four to five times more value than it was capturing in revenue. Every contract reviewed was leaving $450–650 on the table.

The diagnosis came from a newly appointed CFO who, at his first board meeting, asked a question that no one had asked before: "What would we charge if we were pricing this product as a professional service rather than as a software tool?" When he presented the analysis showing that comparable human work cost $600–800, the room went quiet. The board had been congratulating itself on revenue growth. The CFO was showing them that they had been systematically underpricing their way through that growth.

This is the discovery moment for the open-claw problem. It comes when someone in the organization asks not "what are we charging?" but "what could we defensibly charge if we anchored our price to the value we create?" The gap between those two numbers is the measurement of the open claw.

> ***"The right question is not 'how much revenue are we generating?' It is 'how much value are we creating, and what fraction of it are we capturing?' The ratio between those two numbers is the state of your claw."***

Open-Claw Diagnostic — Three Questions

Question	How to answer it	What the answer reveals	Action threshold
1. What value does our AI create per unit of work?	Define the value unit (review, resolution, task). Establish the economic value of the human alternative. Multiply by AI volume.	The total value pool that your AI creates annually — the denominator for your capture rate calculation.	If you cannot answer this question, your measurement infrastructure is missing. Build it before proceeding.
2. What are we	Divide annual revenue by total value created	Your capture rate — the percentage of the value	Capture rate below 10%: significant open claw. Capture

capturing per unit of value created?	(from Q1). Express as a percentage.	you create that converts to revenue. The open claw is 100% minus this number.	rate below 5%: urgent intervention required. Capture rate above 20%: strong commercial position for this product.
3. At what rate is the claw opening or closing?	Compare this quarter's capture rate to last quarter's. Is the gap widening or narrowing?	The trajectory — whether the commercial model is keeping pace with AI capability improvement or falling further behind.	Widening gap despite revenue growth: immediate commercial model intervention. Stable gap: evolutionary commercial development. Narrowing gap: continue current trajectory.

Chapter One — The Essentials

› The discovery moment: comparing the value the AI creates to the revenue it generates. The gap is the open claw.

› The legal AI example: $150 charged per review vs $600–800 value delivered = 20–25% capture rate. Four to five times the current revenue was commercially defensible.

› Two numbers define the claw: value created per unit (determined by the human alternative cost) and revenue captured per unit (the billing rate).

› The diagnostic has three questions: how much value per unit, what fraction is captured, and is the gap widening or narrowing?

› A widening claw during revenue growth is the most dangerous signal: the company is growing while simultaneously losing more value per dollar of growth.

CHAPTER TWO

The Four Forces That Open the Claw

Pricing lag · measurement gap · contract lock · race to zero. How each force operates and compounds.

The open claw does not open randomly. It opens because of four specific forces that are systematically present in AI commercial operations, each independently capable of creating a gap, all four typically operating simultaneously.

The Pricing Lag Force: AI capability improvements arrive continuously — with each new model version, each product iteration, each new training cycle. Commercial pricing adjustments arrive at contract renewal cycles — annually or every three years for most enterprise contracts. This creates a systematic lag: by the time a pricing adjustment can be negotiated, the AI's capability has advanced further. The lag is not laziness or inattention. It is structural: enterprise procurement processes simply do not move at the speed of AI development.

The measurement gap force is the second driver: companies often cannot measure the value their AI creates because they have not built the infrastructure to do so. An AI that improves a customer's operations is creating economic value — time saved, errors avoided, revenue accelerated, costs reduced. But if the measurement system does not capture these improvements, the value is invisible. Invisible value cannot be priced. The inability to measure becomes the inability to capture.

This is the force that most directly responds to commercial intervention. The measurement gap is not technological — the data for measuring AI value typically exists in the customer's systems. It is architectural: the measurement tools, the data pipelines, and the reporting infrastructure must be deliberately built. The company that invests in measurement infrastructure is building the prerequisite for closing its claw. The company that does not is permanently limited in its ability to justify the prices its AI capabilities could command.

The contract lock force is the third driver, and the most commercially frustrating. When an enterprise signs a three-year contract with AI pricing that reflects the technology's current capabilities, they acquire a contractual right to receive the AI's services at those prices for the duration of the contract — regardless of how much better the AI becomes during those three years. A company that signed an AI automation contract in 2022 and is receiving AI capabilities in 2025 that are three times more powerful is paying 2022 prices for 2025 capabilities. The vendor's AI investment has dramatically outpaced the commercial model's ability to capture the returns.

The race-to-zero force is the fourth driver, and the one most external to any individual company's decisions. Competitive pressure to acquire customers often leads AI companies to underprice relative to value during their growth phase — to win market share by being cheaper than the value they create would justify. This is a reasonable strategy for acquisition, but it creates a systematic open-claw problem: the entire market becomes anchored to pricing that was set to acquire customers rather than to capture value.

All four forces operate simultaneously in most AI commercial operations, which is why the open-claw effect is nearly universal. The question is not which companies have an open claw. It is which companies have a strategy to close theirs.

The Four Claw-Opening Forces — Reference

Force	Mechanism	Typical magnitude	Detection signal	Specific closer
Pricing lag	AI capability improves continuously; contract pricing adjusts at renewal cycle. Gap widens between renewal events.	10–30% annual capability improvement; 0–5% annual pricing improvement in existing contracts	Current product capability significantly exceeds the capability level priced in existing contracts	Elastic contracts with capability adjustment provisions; proactive renewal with value-based repricing
Measurement gap	AI creates value that is not measured and therefore cannot be priced. Invisible value stays uncaptured.	Can be 80%+ of total value if measurement infrastructure is absent	Unable to articulate specific dollar value of AI impact; commercial conversations are feature-based, not value-based	Measurement infrastructure investment as prerequisite to any commercial model evolution
Contract lock	Long-term contracts	5–20% annual	NRR below potential	Elastic contract design for new

	signed at historical capability levels lock pricing. The 3-year contract signed at 2022 pricing pays 2022 prices for 2025 capabilities.	pricing gap for each year of contract term remaining	despite AI capability improvement; customers getting dramatically more value than at contract signing	contracts; proactive capability communication for locked contracts; early renewal incentives
Race to zero	Customer acquisition pricing set below value creates market anchors that constrain future pricing. The low price becomes the perceived market rate.	Varies; can be 50–80% of value in early market-building pricing	Price benchmarks set during land-and-expand phase constrain expansion pricing	Segment-specific pricing; tier introduction above existing pricing; outcome pricing as alternative to usage pricing for new customers

THE COMPOUND EFFECT

All four forces operate simultaneously — which is why the claw opens faster than any single force would suggest

The four forces are not independent. When pricing lag keeps existing contracts below value, the measurement gap prevents new contracts from being priced at value, the contract lock prevents renegotiation of below-value contracts, and the race-to-zero creates market anchors that constrain pricing for all customers — all four forces reinforce the same direction: the gap widens. The company that addresses only one force closes only a fraction of the claw. Effective claw closure addresses all four forces with coordinated strategy.

Chapter Two — The Essentials

› Four forces open the claw: pricing lag (contracts vs capability), measurement gap (invisible value), contract lock (frozen pricing), race-to-zero (acquisition pricing anchors).

› Each force is independently capable of creating a significant gap; all four operating simultaneously creates the compound open-claw effect.

› The pricing lag force is addressed by elastic contracts and capability-adjustment provisions in new contracts.

› The measurement gap force is addressed by measurement infrastructure investment — the prerequisite for every other closing strategy.

› The race-to-zero force is addressed by segment-specific pricing and the introduction of outcome pricing as an alternative to usage pricing anchored at acquisition rates.

CHAPTER THREE

The Claw Across Five Layers: Where the Gap Is Widest

How the open-claw effect manifests differently at compute, model, token, agent, and outcome layers.

The open-claw effect manifests differently at each layer of the AI economy, because the nature of the capability advance, the measurement infrastructure requirements, and the commercial model evolution differ by layer.

At the compute layer, the open claw manifests as efficiency advancement without proportional pricing. AI hardware — GPUs, TPUs, custom AI accelerators — has improved at a rate that dramatically exceeds traditional Moore's Law trajectories. The cost of running a trillion-parameter model inference has fallen by approximately 10× in two years. Compute providers who signed multi-year GPU rental agreements in 2022 at 2022 GPU economics are receiving dramatically more compute value per dollar than they were paying for. The open claw at the compute layer is mostly closed by market dynamics — spot pricing for AI compute fluctuates with supply and demand, and enterprise compute contracts include price adjustment provisions. But for companies that locked in long-term GPU commitments before the efficiency explosion, the compute claw is real.

At the model layer, the open claw manifests as capability per dollar improvement. Claude 3.5 Sonnet in late 2024 offers capabilities dramatically superior to GPT-3.5 in early 2023 — at a comparable or lower price per million tokens. Companies that built applications on earlier, more expensive model versions and charged customers prices that reflected those earlier costs are now delivering dramatically better performance at dramatically lower cost, without a corresponding increase in what they charge. The per-token cost has fallen; the per-interaction value has risen; the commercial model has not updated to capture the difference.

At the token layer, the open claw manifests as context window and reasoning improvements that create value without price adjustment. When GPT-4 expanded its context window from 8K to 128K tokens in 2023, the value of a single inference request for document-intensive tasks multiplied dramatically — the same token price now purchased the ability to analyze an entire legal agreement rather than a few pages at a time. The price per million tokens did not change. The value delivered per million tokens increased substantially.

At the agent layer, the open claw manifests most dramatically, because agent capabilities are advancing the fastest and the commercial models for agent work are the least mature. A company that deployed AI agents in 2022 to handle simple, defined tasks and priced them at $X per task is now operating agents that can handle dramatically more complex tasks — tasks that would cost 5–10× more if performed by human professionals — at the same $X per task pricing. The agent's capability has compounded; the commercial model has not.

At the outcome layer, the open claw manifests as measurement underinvestment. Companies that are positioned to charge for outcomes are often unable to because they have not built the verification and measurement infrastructure that makes outcome billing credible. The AI creates outcomes — verifiable, economically valuable outcomes — that the measurement gap prevents from being captured in commercial terms.

Understanding which layer of the AI economy creates the most significant open-claw exposure is the first step in a layer-specific closing strategy.

Open-Claw Gap by AI Economy Layer

Layer	How capability improves	How revenue captures it (or doesn't)	Gap mechanism	Widest gap at
Compute	GPU/TPU efficiency improving 2–3× per year; same inference cost buys dramatically more compute	Multi-year reservations lock pricing; spot pricing partially elastic but lags hardware efficiency	Reserved capacity contracts priced on 2022 GPU economics receiving 2025 GPU efficiency	Enterprise customers who locked long-term GPU commitments before efficiency explosion
Model	Capability per token improving dramatically; GPT-4o vs GPT-3.5 at comparable price	Per-token prices declining as competition intensifies and efficiency improves; value per token rising	Value per token rising while price per token falls — the efficiency improvement benefits the customer, not the vendor	High-value use cases where the same token count now produces dramatically more useful outputs
Token	Context windows expanding; reasoning quality improving per token	Billed by token count regardless of capability improvement — more tokens in context window = more billing, but more value per token	Token count is a proxy for consumption, not for value; as reasoning quality improves, the same token count creates more value	Complex analytical use cases where large context windows enable qualitatively different outputs
Agent	Task completion quality and reliability improving rapidly; 2024 agents do what 2022 agents could not	Per-task pricing not yet widespread; most agent products still on subscription or per-token	The agent's ability to complete complex tasks grows while per-task (or per-token) pricing stays constant	Enterprise agent deployments where the scope of addressable tasks has expanded dramatically since contract signing

Outcome	Outcome delivery rate improving as AI quality improves; more outcomes per deployment	Outcome pricing not yet widespread; most companies still on consumption or subscription	AI improving outcome delivery rate while companies still not pricing on outcomes — the value improvement is entirely customer benefit	Healthcare, legal, and financial AI where outcome quality improvement is measurable but commercially uncaptured

Chapter Three — The Essentials

› The claw manifests differently at each AI economy layer — compute (efficiency gain), model (capability per token), token (context/reasoning quality), agent (task complexity), outcome (delivery rate).

› The agent and outcome layers have the widest current gap — agent capabilities are improving fastest, and outcome pricing is least mature.

› The compute layer claw is partially self-closing through market dynamics (spot pricing, contract renegotiation).

› The model layer claw is structurally open in per-token pricing — efficiency improvement benefits the customer, not the vendor.

› The most actionable layer for claw closure is the agent layer — where capability improvement is dramatic, the commercial model is least mature, and the opportunity is largest.

CHAPTER FOUR

The Claw as Enabler: When Closing Faster Creates Compounding Advantage

Why the company that closes its claw first wins disproportionately. The compounding mechanics.

The open claw is universally discussed as a problem. It is also an opportunity — and understanding the opportunity dimension is essential to developing the right strategic response.

The opportunity is this: the company that closes its claw faster than its competitors captures a compounding commercial advantage. When an AI capability creates value, there is a period during which that value is either captured (by the vendor with the pricing model to capture it) or not captured (it accrues to the customer as unpriced value). The first company in a market to build the measurement infrastructure, the commercial model, and the organizational capability to price that value creates a revenue stream that competitors who are still leaving the value uncaptured do not have.

This advantage compounds because revenue funds further AI investment, which creates more capability, which creates more value, which — if captured — creates more revenue. The company with the highest capture rate is the company that can reinvest the most aggressively in further capability improvement. The company with the lowest capture rate is the company whose AI investment yields the lowest commercial return, which constrains its ability to fund the next capability advance.

Microsoft's commercial success with GitHub Copilot illustrates the compounding advantage of being the first mover to close the claw in AI-assisted software development. When GitHub Copilot launched in 2022 at $10/user/month ($19/month for individuals, $19/user/month for business), it created the first commercial capture mechanism for AI coding assistance value. The value the product created was significantly higher than $10/user/month — studies documented 30–50% productivity improvements, worth hundreds of dollars per developer per month in time savings. But $10 was the market-clearing price at which developer adoption occurred at scale.

The capture at $10/user/month funded GitHub's continued development of Copilot, which improved the product, which allowed GitHub to introduce Copilot Business at $19/user/month and Copilot Enterprise at $39/user/month as capabilities expanded. Each commercial tier captured more of the growing value. The claw closed progressively as commercial tiers were added.

Microsoft's compounding advantage: the revenue from Copilot's 1.3M+ paid users (approximately $30M/month at blended average pricing) is funding further Copilot development, which is building the capability that will justify the next pricing tier. The first mover who closed the claw first is compounding fastest.

The lesson: the open-claw framework is not only about recovering value that is currently being missed. It is about building the commercial advantage that enables faster compounding than competitors who are closing their claws more slowly.

The Compounding Advantage of First-Mover Claw Closure

Stage	First mover (closes claw at T=0)	Late mover (closes claw at T+2 years)	Gap at T+4 years
Year 1	Closes claw to 15% capture rate. Revenue = $15M on $100M value created.	Capture rate = 5%. Revenue = $5M on $100M value created.	First mover has $10M more annual revenue to reinvest in AI.
Year 2	Reinvests $10M in AI development. AI improves; creates $150M value. Capture rate grows to 18%. Revenue = $27M.	No reinvestment differential. AI improves at base rate; creates $130M value. Capture rate stays 5%. Revenue = $6.5M.	Revenue gap: $20.5M. Capability gap beginning to compound.
Year 3	Continued reinvestment differential builds capability lead. Creates $225M value at 20% capture = $45M revenue.	Competitor now investing to close claw (too late). Creates $170M value at 8% capture = $13.6M revenue.	Revenue gap: $31.4M. First mover has 3× higher revenue funding 3× faster AI development.
Year 4	Capability and commercial lead fully established. Creates $340M value at 22% = $74.8M.	Playing catch-up from significantly behind. Creates $210M value at 12% = $25.2M.	First mover 3× larger by revenue, compounding through reinvestment at 3× rate.

> ***"The first company in a market to close its claw captures the revenue that funds the next capability generation. That capability generates the next value surplus. The***

> ***closed claw enables the compounding that makes the first mover permanently advantaged."***

Chapter Four — The Essentials

- › The claw as enabler: the company that closes its claw first captures the revenue that funds faster AI development — creating compounding advantage.
- › The compounding math: at T+4 years, a first mover who closed at T=0 has 3× the revenue of a competitor who closed at T+2.
- › Microsoft/GitHub Copilot demonstrates the compounding advantage: each tier launch captures more value that funds the next tier's development.
- › The first-mover advantage in claw closure is not just commercial — it is also the funding mechanism for the capability lead that makes the commercial position defensible.
- › The urgency: every quarter of delay costs two quarters of compounding advantage — the quarter lost and the quarter of compounding that the quarter's revenue would have funded.

PART TWO

Measuring the Claw

Quantification methodology. Building the infrastructure. The board-level dashboard.

CHAPTER FIVE

Quantifying Your Capability-Capture Gap

The measurement methodology. How to calculate the dollar value of uncaptured capability.

Quantifying the open-claw gap requires answering two questions: what value is the AI actually creating, and what is the company actually capturing? The gap between these two numbers is the claw.

Both questions are harder to answer than they appear. The "value created" question requires measurement infrastructure that most companies have not yet built. The "revenue captured" question is straightforward — it is the revenue line — but the calculation of the gap requires that both numbers be expressed in comparable units and comparable timeframes.

The measurement methodology for the open-claw gap:

Step one is defining the value unit. What is the natural unit of value that the AI creates in this specific deployment? For a contract review AI, the value unit is a reviewed contract — the economic value of having a contract reviewed by the AI versus the cost of the human alternative. For a customer service AI, the value unit is a resolved inquiry — the cost of AI resolution versus the cost of human resolution. For a code generation AI, the value unit is a code contribution — the time saved per developer versus their fully-loaded hourly cost. The value unit must be specific, measurable, and economically meaningful.

Step two is establishing the baseline value of one unit. What does the alternative cost? The customer service AI that resolves a ticket in 3 minutes is displacing a human agent who would resolve the same ticket in 12 minutes at a fully-loaded cost of $35/hour. The baseline value of one AI resolution is the human cost of the alternative: (12 minutes / 60 minutes) × $35 = $7 in human resolution cost displaced. This is not the value the customer captures from having the inquiry resolved — it is the cost of the human alternative that the AI displaces. If the customer values the inquiry resolution at more than $7 (which they do, because they would not have staffed the contact center otherwise), the AI value calculation is even more conservative.

Step three is establishing the current capture rate. What is the company charging per unit, expressed in the same terms as the baseline value? If the customer service AI charges $0.99 per resolution and the baseline value is $7 per resolution, the capture rate is 14%. The open-claw gap is 86%.

Step four is calculating the total gap. Annual AI interactions × baseline value per interaction − annual revenue from those interactions = total open-claw gap in dollar terms. For the customer service AI: 10 million resolutions per year × $7 value = $70M in total value created. Revenue captured: 10 million × $0.99 = $9.9M. Total open-claw gap: $60.1M per year — value created but not captured.

Step five is analyzing the gap composition. Not all of the open-claw gap is capturable — the customer must retain enough value from the AI's work to justify using it. If the customer captures none of the value (the vendor takes it all), there is no commercial relationship. The commercially appropriate capture target — the percentage of value that the vendor should aim to capture — depends on the competitive dynamics, the customer's alternatives, and the strength of the vendor's commercial position. A range of 10–30% is typical for professional AI services; higher capture rates are defensible in monopolistic or highly differentiated market positions.

For the customer service AI, at a 20% target capture rate, the commercially appropriate price is $7 × 20% = $1.40 per resolution. Current price is $0.99. The actionable open-claw gap — the revenue increase achievable by pricing toward the target capture rate — is ($1.40 − $0.99) × 10 million = $4.1M per year. Not the entire $60M gap (most of which appropriately accrues to the customer), but a real and significant commercial opportunity.

Open-Claw Quantification — Five-Step Methodology

Step	Activity	Example (customer service AI)	Output
1. Define the value unit	Identify the natural unit of value: what does the AI accomplish that can be measured and	One resolved customer service ticket — the AI closes the inquiry without human escalation.	Value unit: ticket resolution. Measurement: ticket closure without escalation AND not reopened within 24h.

	priced?		
2. Establish baseline value	What does the human alternative cost? This is the economic value of one AI value unit.	Human agent resolution: 12 minutes × ($35/hour ÷ 60 minutes) = $7 per ticket resolved.	Baseline value: $7 per resolved ticket.
3. Establish current capture rate	What is the vendor charging per unit? Divide by baseline value.	Intercom charges $0.99 per resolution. $0.99 ÷ $7 = 14.1% capture rate.	Current capture rate: 14.1%. Open claw: 85.9% of value uncaptured.
4. Calculate total gap	Annual volume × baseline value − annual revenue = total open-claw gap.	10M resolutions × $7 = $70M value created. Revenue: $9.9M. Gap: $60.1M.	Total gap: $60.1M/year. This is the maximum additional value available — not the commercially appropriate target.
5. Identify actionable gap	Target capture rate × baseline value per unit − current price = actionable price increase per unit. Multiply by volume.	Target: 20% capture. Target price: $7 × 20% = $1.40. Actionable increase: $0.41/resolution × 10M = $4.1M/year.	Actionable gap: $4.1M/year — the revenue increase achievable by moving to target capture rate.

THE CAPTURE RATE TARGET

10–30% is the typical commercially appropriate capture range — not 100%

The open-claw gap is not entirely capturable. The customer must retain enough of the value created to have a strong economic incentive to use the AI. If the vendor captures 100% of the value, the customer has zero net benefit — the AI displaces the human cost exactly and provides nothing to the customer beyond the status quo. The commercially appropriate target is 10–30% of value captured, leaving 70–90% for the customer as net benefit. This range gives the customer a compelling ROI (3–10× value vs price) while giving the vendor pricing power that reflects the value created. At the target capture rate, 'why would I pay this?' becomes 'why would I not pay this?'

Chapter Five — The Essentials

› The five-step methodology: define the value unit → establish baseline value → establish current capture rate → calculate total gap → identify actionable gap.

› The total gap is not the actionable gap: most of the value appropriately accrues to the customer as net benefit.

› The target capture rate (10–30%) determines the actionable gap — the revenue increase achievable by moving from current to target pricing.

› The human alternative cost is the most defensible baseline for value measurement: it is independent of AI capability claims and verifiable by both parties.

› Every AI product can be measured using this framework — the specific value unit, baseline cost, and volume metrics vary; the methodology is universal.

CHAPTER SIX

Building the Value Measurement Infrastructure

The systems, processes, and data required to measure AI value before you can price it.

Building the value measurement infrastructure is the prerequisite for every commercial strategy in this book. Without measurement, the value anchor for pricing does not exist. Without measurement, the constraint analysis for variable consideration cannot be performed. Without measurement, the outcome-based contract cannot be defended. Without measurement, the gain-share calculation has no credible basis.

The value measurement infrastructure for AI products has four components:

Baseline measurement establishes the pre-AI state of the metric the AI is affecting. This requires either a historical baseline (the metric's value before the AI was deployed) or a controlled baseline (a comparison group not using the AI). The baseline is the reference point for calculating improvement — without it, attribution of the improvement to the AI is contested. The baseline must be established before deployment, not retrospectively, because retrospective baselines are always disputed.

Outcome tracking captures the metric's value after AI deployment, with sufficient granularity to attribute changes to the AI's activity rather than to other factors. This requires instrumented integrations: the AI's activity must be logged alongside the customer's outcome metrics in a way that enables causal analysis. For a contract review AI, this means logging which contracts were AI-reviewed and subsequently tracking the customer's legal department's cycle time, dispute rate, and attorney time per contract — metrics that reflect whether the AI review improved outcomes.

Attribution analysis separates the AI's contribution from the other factors affecting the outcome metric. This is the most technically demanding component of the measurement infrastructure, because outcome improvements have multiple causes. Statistical approaches (difference-in-differences analysis, A/B testing between AI-assisted and non-AI-assisted workflows) are the most credible for attribution, but they require the organizational access and data quality to implement.

Value translation converts the measured improvement into dollar terms. A contract review AI that reduces average contract cycle time from 21 days to 9 days has created a 12-day improvement. The dollar value of 12 days depends on: the customer's contract value (faster cycles mean faster revenue recognition), the cost of attorney time avoided in the review process, and the risk reduction from more thorough review (fewer unfavorable contract terms accepted because the AI caught them). Each component must be calculated at the customer's specific economics, not at a generic industry average.

The measurement infrastructure is not cheap to build. A comprehensive value measurement program for an enterprise AI product typically requires three to six months of data engineering, instrument development, and customer integration work. It requires ongoing statistical analysis capability. And it requires customer cooperation — the customer must provide access to the outcome data that the measurement requires.

This investment is why most AI companies have not built it yet. It feels like a non-revenue activity. It feels like research. It is, in fact, the highest-return commercial investment an AI company can make — because every dollar spent building measurement infrastructure creates the foundation for pricing that captures multiples of that dollar in additional revenue.

Value Measurement Infrastructure — Four Components				
Component	**What it captures**	**Implementation requirement**	**Time to build**	**Commercial prerequisite for**
Baseline measurement	Pre-AI state of outcome metrics — the reference point for calculating improvement	Historical data extraction + statistical baseline setting + customer access agreement. Must be done BEFORE deployment.	4–8 weeks per use case	Attribution analysis · Outcome anchoring · Gain-share models
Outcome tracking	Post-AI state of outcome metrics with sufficient granularity for attribution analysis	Integration with customer's system of record (CRM, ticketing system, ERP). Bidirectional data flow. Automated metric calculation.	3–6 months for full integration	Attribution analysis · Outcome invoicing · Performance reporting
Attribution analysis	Separation of AI's contribution from other factors affecting the outcome metric	Statistical methodology (DID, A/B test, regression) + sufficient observation volume (typically 500+ interactions). Requires data science capability.	6–12 months for validated methodology	Outcome anchoring · Gain-share · Capability adjustment provision
Value translation	Dollar value of the measured improvement at the customer's specific economics	Customer-specific economic model: cost per unit, volume, conversion value. Requires customer data sharing and economic modeling capability.	2–4 weeks once attribution is validated	Value-based pricing · Value anchor for sales conversations · ROI reporting

Measurement Infrastructure by AI Product Category				
Product category	**Value metric**	**Baseline source**	**Attribution challenge**	**Dollar translation**
Productivity AI (coding, writing, knowledge work)	Time saved per task category	Time tracking data or sampling study before	Attribution is challenging — multiple factors affect productivity;	Fully-loaded cost per hour for the role × hours saved per period

		deployment	A/B deployment helps	
Process automation AI (workflow, data processing)	Cycle time, error rate, cost per transaction	Process metrics from existing operational systems (cycle time logs, error records)	Attribution is relatively clean — process metrics are directly affected by AI deployment	(Pre-AI cycle time − Post-AI cycle time) × volume × cost per unit time OR error rate reduction × average error cost
Decision support AI (forecasting, analytics, recommendations)	Forecast accuracy, decision quality, downstream outcome improvements	Historical forecast records, decision logs, outcome tracking	Complex — must separate AI contribution from improved data, market conditions, team learning	Forecast error reduction × cost of forecast errors (e.g., inventory cost per unit of forecast error, deal slippage from forecast inaccuracy)
Customer-facing AI (support, sales, engagement)	Resolution rate, satisfaction score, conversion rate, ticket deflection	Historical contact center metrics, NPS scores, conversion tracking	Attribution requires controlled comparison (AI-handled vs human-handled for comparable inquiry types)	Human agent cost per resolution × resolution volume + satisfaction improvement × LTV impact
Professional AI (legal, medical, financial)	Task completion time, quality score (expert review), compliance accuracy	Time tracking from billing systems, quality review results, compliance	Attribution challenging in professional services; controlled quality audits are	Professional billing rate × time saved per task OR (Human error rate − AI error rate) × average

		audit records	most credible	cost of errors in domain

⚠ Measurement Infrastructure Takes 12–18 Months — Start Now

The most common claw-closing failure pattern: an AI company recognizes the open-claw problem, decides to implement outcome-based pricing, attempts to implement it without measurement infrastructure, and fails because the evidence does not exist to defend the outcomes being billed. The attempt is abandoned. The company returns to subscription or per-token pricing. The claw continues to widen. The correct sequence is: measure first, price second. Not the other way around. If your measurement infrastructure is not being built today, your commercial model evolution is 12–18 months away regardless of when you decide to start.

Chapter Six — The Essentials

› Measurement infrastructure has four components: baseline measurement, outcome tracking, attribution analysis, value translation.

› Attribution analysis is the most technically demanding component — it requires statistical methodology and typically 6–12 months of validated operation.

› Baseline measurement must be established before AI deployment — retroactive baselines are always disputed.

› The measurement infrastructure timeline is 12–18 months for a complete, defensible system — this is not optional or compressible.

› Start building measurement infrastructure now, regardless of when commercial model evolution is planned — the infrastructure must precede the commercial model change, not follow it.

CHAPTER SEVEN

The Claw Dashboard: Tracking Capability vs Capture Over Time

How to report on the claw gap to your board. The metrics that matter.

The Claw Dashboard is the management reporting instrument that makes the open-claw gap visible to the board and the executive team. Without it, the claw opens silently — the revenue line is growing, which masks the growing gap between what is being charged and what could be charged.

The Claw Dashboard has five panels:

Panel 1: The Claw Gap Over Time. A time series chart showing, for each quarter, the total value created by the AI (measured in the units established by the value measurement infrastructure) and the total revenue captured from that value. The gap between the two lines is the open claw, visualized. This panel should make the board viscerally aware of the magnitude and trajectory of the gap. If the gap is widening every quarter (which it typically is when AI capability is advancing faster than commercial model evolution), this panel is the most important governance signal the board receives.

Panel 2: Capture Rate by Product and Customer Segment. The percentage of value created that is captured in revenue, broken down by product line and customer segment. This panel identifies where the claw is widest — which products and which customer segments have the largest gap between value created and value captured. It also identifies early wins: the segments where pricing is already close to value-based anchors.

Panel 3: Gap Composition Analysis. The breakdown of the total open-claw gap by the four forces that drive it: pricing lag (revenue lost to below-market prices in existing contracts), measurement gap (value created but not yet quantified and therefore not priced), contract lock (revenue foregone due to long-term contracts signed at below-current pricing), and race-to-zero legacy (systematic underpricing inherited from customer acquisition pricing). This decomposition identifies the specific interventions required — each gap driver requires a different closing strategy.

Panel 4: Closing Rate and Trajectory. The rate at which the claw is being closed — the quarterly improvement in capture rate — and the projection of future capture rate at the current closing pace. This panel tells the board whether the company's commercial model evolution is keeping pace with its AI capability evolution. If the closing rate is slower than the claw opening rate (which it typically is in high-growth AI companies), the gap is widening even as the company is making progress on its commercial model.

Panel 5: Competitive Benchmark. A comparison of the company's capture rate against the known or estimated capture rates of comparable companies in the same market. This panel contextualizes the gap: a 15% capture rate may be the industry standard, in which case the commercial opportunity is smaller than if the industry average is 25%. It also identifies the competitive opportunity: the company that achieves a 25% capture rate in a market where the average is 15% has a revenue advantage that compounds through the AI investment flywheel.

The Claw Dashboard — Five Panels

Panel	What it shows	Key metric	Board threshold for concern	Reporting cadence
1. Claw Gap Over Time	Time series: value created vs revenue captured per quarter. The gap between the two lines is the open claw.	Quarterly gap in dollar terms. Gap growth rate.	Gap growing at > 10% per quarter indicates commercial model evolution is not keeping pace with capability improvement.	Quarterly
2. Capture Rate by Segment	Capture rate (revenue/value) broken down by product line and customer segment.	Capture rate % by segment. Segment trend.	Any segment below 5% capture rate warrants immediate commercial model review.	Quarterly
3. Gap Composition Analysis	Breakdown of gap by the four forces: pricing lag, measurement gap, contract	Dollar value attributed to each force. Force trend.	Measurement gap component above 50% indicates measurement	Semi-annual

	lock, race-to-zero.		infrastructure is the primary bottleneck.	
4. Closing Rate and Trajectory	Rate of capture rate improvement. Projection of future capture rate at current closing pace.	Quarterly capture rate improvement. Months to target capture rate.	Negative closing rate (claw widening despite commercial effort) requires strategic intervention.	Quarterly
5. Competitive Benchmark	Comparison of company's capture rate against estimated competitor capture rates.	Relative capture rate. Competitive gap.	Capture rate below estimated industry average warrants analysis of competitive commercial model differences.	Annual

BOARD GOVERNANCE

The Claw Dashboard belongs in every board pack alongside ARR, NRR, and gross margin

Most AI company boards are managing their AI investments by looking at the revenue side of the equation — how much revenue is being generated. The Claw Dashboard adds the value side of the equation — how much value is being created and what fraction is being captured. Without the value side, the board cannot assess whether the company's AI investment is generating appropriate returns or whether the open claw is systematically undermining the commercial model. A board that does not see the claw cannot govern for its closure. Make it visible. Make it a quarterly governance metric.

Chapter Seven — The Essentials

› The Claw Dashboard has five panels: gap over time, capture rate by segment, gap composition, closing rate, and competitive benchmark.

› The most important single metric: capture rate trend. Declining capture rate despite revenue growth is the signature of a widening claw.

- › The gap composition panel identifies which of the four forces is primary — each requires a different closing strategy.
- › The competitive benchmark panel contextualizes the gap: below industry average capture rate represents both a vulnerability and an opportunity.
- › Include the Claw Dashboard in the quarterly board pack — gap visibility is the governance prerequisite for claw closure.

PART THREE

The Four Closers

Outcome anchoring · Elastic contracts · Measurement first · Moat building.

CHAPTER EIGHT

Closer 1 — Outcome Anchoring: Price What the AI Does, Not What It Is

The strategy of anchoring price to measurable business outcomes. Implementation mechanics.

Outcome anchoring is the commercial strategy of pricing what the AI accomplishes rather than what the AI is. It is the most direct closer for the open-claw gap because it connects the price to the value measurement directly — the price is a function of the outcome, and the outcome is the thing being measured.

The mechanism is straightforward: define a specific business outcome that the AI creates, establish the measurement methodology for verifying that the outcome occurred, and set the price as a fraction of the economic value of that outcome. When the AI improves (as it will), the outcome is delivered more reliably, more quickly, or at higher quality — but the economic value of the outcome is unchanged. The price per outcome remains appropriate. The claw stays closed.

Compare this to per-token or per-seat pricing. When the AI improves under per-token pricing, the same outcome is achieved with fewer tokens — the customer gets more value for the same price. This is wonderful for the customer but creates a claw-opening dynamic for the vendor: the AI's improvement directly translates into a lower revenue-per-outcome as efficiency increases. The per-token model is structurally claw-opening in a world of AI capability improvement.

Under per-outcome pricing, AI improvement creates more reliable outcome delivery, faster outcome delivery, or higher-quality outcome delivery — but the price per outcome is stable. The claw stays closed because the value being priced (the outcome) does not diminish as the AI improves. It may actually increase: a more reliable outcome (the AI gets it right 99% of the time instead of 95%) is worth more to the customer, justifying a gradual price increase aligned with reliability improvement.

Specific implementation of outcome anchoring across different AI applications:

Intercom's resolution-based pricing at $0.99 per resolved customer service ticket is outcome anchoring in its most mature commercial form. Intercom is not charging per token, per conversation, or per user license. It is charging per resolved ticket — a specific, verifiable outcome with a defined economic value (the human agent resolution cost it displaces). As Intercom's AI improves and resolves more complex tickets more reliably, the price per resolution remains $0.99 — but the volume of resolvable tickets increases, generating more revenue from the same pricing structure.

Harvey AI's per-review pricing for legal documents is outcome anchoring applied to professional services delivery. Harvey charges per contract review delivered — a specific work product with a defined economic value (the junior associate time it displaces). As Harvey's AI improves and produces higher-quality reviews, the value of each review delivered increases, but the price per review increases proportionally with quality tier.

The outcome anchor (a review delivered) stays consistent; the quality tiers (Standard, Professional, Expert) allow price to rise with capability improvement.

Workato, the enterprise automation platform, has moved portions of its commercial model toward outcome-based pricing tied to workflow automation ROI. Enterprise customers who can demonstrate specific process improvement metrics (cycle time reduction, error rate improvement, cost per transaction reduction) have access to pricing structures that are tied to the measured improvement rather than to the number of users or workflows automated. This is outcome anchoring in the workflow automation context.

The two requirements for successful outcome anchoring are precision and verifiability. The outcome must be precisely defined — specific enough that both parties agree unambiguously on whether it occurred. And it must be verifiable — the determination of whether the outcome occurred must be based on data from the customer's systems rather than from the vendor's AI system (which has an inherent conflict of interest in self-reporting). Harvey's model of independent attorney verification of review quality is an example of verifiability design: neither party controls the verification, which makes the outcome pricing credible.

CLOSER 1: OUTCOME ANCHORING

Connect the price to the verified outcome. When the AI improves, the outcome value is maintained. The claw stays closed.

Core logic	Outcome anchoring closes the claw structurally by pricing the output (the outcome) rather than the input (the tokens, the API calls, the time). When the AI improves and delivers the same outcome more efficiently, the price per outcome is stable — the efficiency improvement does not erode revenue as it does under input pricing.
Implementation sequence	(1) Define the outcome with precision: binary, measurable, attributable, economically meaningful, completable (see Book 5 Chapter 9). (2) Build verification infrastructure. (3) Calculate baseline value. (4) Set price at target capture rate. (5) Track and report on outcome delivery rate.
The anti-claw-opening property	Under per-token pricing: AI improves → same outcome achieved with fewer tokens → customer gets same value for lower cost → vendor revenue per outcome declines. Under per-outcome pricing: AI improves → same outcome

	achieved more reliably → customer satisfaction increases → outcome volume may grow → vendor revenue stable or growing per outcome delivered.
The quality premium	As the AI improves, outcome quality improves. Better quality outcomes can be priced at a premium: the Standard tier ($150/review, 85% recall) can coexist with a Professional tier ($450/review, 95% recall) that captures the value of the capability improvement in a natural commercial structure.
When to introduce	Outcome pricing can be introduced selectively — for new customer segments, new use cases, or new contract vintages — without requiring migration of the entire installed base. The selective introduction allows measurement validation and commercial learning before broader deployment.

CASE STUDY: INTERCOM *Resolution Pricing — The Gold Standard of Outcome Anchoring*	
The commercial structure	Intercom charges $0.99 per resolved customer service inquiry — defined as: the customer's issue was resolved by the AI without human involvement, and the conversation was not reopened within 24 hours. This is outcome pricing in its purest form.
The claw-resistant property	Under Intercom's $0.99/resolution pricing, AI improvement does not open the claw. As Intercom's AI becomes more capable and resolves more complex inquiries reliably, two things happen: (1) the volume of resolvable inquiries grows — more complex issues that previously required human escalation can now be AI-resolved; and (2) the value per resolved inquiry increases as more valuable issues are resolved. Revenue grows with capability improvement, not despite it.
The capture rate evolution	At $0.99/resolution and a human agent cost of $7/resolution, Intercom captures 14.1%. As Intercom's AI resolves more complex inquiries (worth $15–20 in human agent time for complex issues), the same $0.99 price represents an even lower capture rate for those inquiries. The claw is still partially open — Intercom could price more aggressively — but the outcome pricing structure prevents the claw from opening further as capability improves.
The commercial contrast	Compare to a per-conversation pricing model at $0.50/conversation: as AI becomes more efficient, the same resolution requires fewer conversation turns, generating less revenue per resolution. The per-conversation model is structurally claw-opening. The per-resolution model is structurally claw-stable.
The expansion mechanic	Intercom's outcome pricing creates a natural commercial expansion mechanic: as Intercom's AI resolves a higher percentage of incoming inquiries (from 40% to 60% to 80% as capability improves), Intercom's revenue from the same customer grows proportionally — without any price increase, just from the expansion of the AI's scope.

CASE STUDY: WORKATO *Process Automation — Outcome Anchoring for Workflow ROI*	
The evolution	Workato has evolved from pure subscription pricing (which does not capture the value of workflow automation outcomes) toward outcome-correlated pricing structures that tie commercial terms to measurable workflow ROI.
The outcome definition challenge	Process automation outcomes are more complex to define precisely than customer service resolutions. A 'successful workflow automation' is not binary — it has degrees of quality, reliability, and scope. Workato's approach: define outcomes in terms of the process metrics that matter to the customer (invoice cycle time, error rate, cost per processed record) rather than in terms of AI activity (workflows run, API calls made).
The measurement infrastructure	Workato's ROI measurement infrastructure connects the commercial model to the customer's operational metrics: cycle time before vs after workflow automation, error rate before vs after, cost per transaction before vs after. This measurement is not trivial to implement — it requires integration with the customer's ERP and operational systems — but it creates the evidence base for outcome-correlated pricing.
The commercial implication	Enterprise customers whose workflow automation delivers a 40% reduction in invoice processing cycle time (from 12 days to 7 days) and a 60% reduction in error rate have achieved outcomes with specific economic values. Workato's pricing for these customers can be anchored to those outcomes — and the anchor makes the pricing defensible against competitive alternatives that cannot demonstrate comparable outcome delivery.

Chapter Eight — The Essentials

› Outcome anchoring closes the claw structurally: the price is anchored to the output, not the input, so AI efficiency improvements do not erode revenue.

› Intercom's $0.99/resolution is the commercial standard — specific definition, customer-system verification, natural expansion mechanic as capability improves.

› The anti-claw-opening property: under outcome pricing, AI improvement increases outcome reliability and volume, growing revenue rather than eroding it.

› The quality premium: tier structure (Standard, Professional, Expert) allows price to grow with quality improvement — capturing capability improvement commercially.

› Introduce outcome pricing selectively before deploying broadly — measurement validation and commercial learning precede scale.

CHAPTER NINE

Closer 2 — Elastic Contracts: Building Expansion Into the Agreement

Consumption ratchets · expansion clauses · usage-based upsells · the contract that grows with value.

Elastic contracts are the commercial infrastructure strategy for closing the claw proactively — building the mechanisms for price expansion into every contract at signing, so that AI capability improvement translates automatically into revenue capture rather than requiring a renegotiation cycle.

The open-claw problem is in large part a contract design problem. When enterprise AI contracts are signed with fixed pricing for three years, the pricing model is locked at the capability level of AI at signing. As the AI improves over the three-year term, the capability-capture gap widens — not because the vendor has not improved the product, but because the contract design did not provide for commercial capture of the improvement.

Elastic contracts address this through four specific mechanisms:

Consumption ratchets are provisions that automatically increase the price per unit when consumption exceeds defined thresholds within a billing period. The mechanism: the first 10,000 units in a month are priced at $1.00 per unit; units 10,001 to 25,000 are priced at $0.90; units above 25,000 are priced at $0.80. This is a volume discount structure — but in the context of AI products where consumption increases as the AI is deployed more deeply, the consumption ratchet also captures the expansion value of increasing AI deployment. The customer pays a lower marginal price as volume grows, which encourages adoption; the vendor captures more total revenue as volume grows, which closes the claw on the expansion dimension.

Expansion clauses are provisions that define automatic pricing adjustments when the customer's AI deployment expands beyond the initially contracted scope. An enterprise AI contract for contract review in the legal department includes an expansion clause: "If the customer deploys this product in additional departments or for additional use cases beyond the contracted scope, additional capacity is available at the contracted per-unit rate plus a 15% expansion premium." This clause converts scope expansion from a renegotiation event (which requires sales effort and creates price pressure) to an automatic commercial event (which captures expansion value without renegotiation friction).

Capability adjustment provisions are provisions that define the commercial terms for pricing when the AI's capabilities materially improve during the contract term. These provisions are relatively novel and face the most customer resistance, because they are explicitly about capturing AI improvement value. The mechanism: "If the vendor releases a materially improved capability version during the contract term that the customer wishes to adopt, the parties will negotiate pricing for the improved version in good faith, with a reasonable presumption that pricing will reflect the improvement in value relative to the baseline version." This provision does not guarantee a specific price increase, but it establishes the commercial framework for capturing capability improvement rather than leaving it as an unpriced benefit to the customer.

Usage-based upsells are structured as catalog provisions: the base contract includes defined usage tiers, and additional usage above each tier is automatically available at catalog prices without requiring a new contract. This is the most common elastic contract mechanism and the one that customers most readily accept because it mirrors familiar cloud computing pricing. The commercial effect: as the customer's AI usage grows (which it typically does as the AI improves and the organization deploys it more broadly), revenue grows automatically without requiring sales team intervention.

Designing elastic contracts requires a specific negotiation posture: the vendor must be willing to accept lower per-unit prices at lower volumes in exchange for automatic expansion rights. Customers who accept elastic contract structures often do so because the initial prices are lower than they would be in a flat-rate contract, with the tradeoff that the vendor captures the commercial upside of adoption growth. This is a reasonable trade for both parties: the customer gets lower initial risk; the vendor gets automatic upside from success.

Companies executing elastic contracts successfully:

Snowflake's consumption-based model is the canonical enterprise software elastic contract at scale. Snowflake charges per credit consumed, with credits scaling to compute and storage usage. As customers' data operations grow and they use Snowflake for more analytical workloads, consumption grows automatically and revenue grows with it. Snowflake's NRR above 130% is partially a function of customers expanding their usage — the elastic pricing model captures that expansion without requiring renegotiation.

Datadog's per-host, per-metric pricing model is similarly elastic: as customers deploy more infrastructure that needs monitoring, more hosts are added to the Datadog account and revenue grows proportionally. Datadog's NRR above 120% reflects the same dynamic — the pricing model is designed to capture growth automatically.

The AI-native elastic contract design challenge is ensuring that the contract's consumption definitions capture capability improvement, not just volume growth. A contract that prices per token will capture volume growth but will see the value-per-token increase as AI efficiency improves without a corresponding price increase. Elastic contracts for AI products should price per outcome (not per token) with elastic provisions for volume — capturing both the growth of the customer's AI deployment and the increasing value of each outcome as the AI improves.

CLOSER 2: ELASTIC CONTRACTS

Design every contract for the AI's success. Build the commercial terms to capture adoption growth and capability improvement automatically — without renegotiation.

Core principle	Every new contract should include provisions that allow revenue to grow as AI deployment deepens, as AI capability improves, and as the customer's AI usage

	expands. The contract should grow with the AI — automatically, without requiring a new sales cycle for each expansion.
Consumption ratchets	Price per unit decreases as volume increases (volume discount), but total revenue increases as volume grows. Customer incentive: lower marginal cost for heavy usage. Vendor benefit: automatic revenue growth with usage growth. Critical design: ratchet thresholds must be set at the expected growth points, not at unreachable volumes.
Expansion clauses	Define the terms for expansion beyond the contracted scope before the expansion happens. Avoid making the customer negotiate for the right to expand — make expansion automatic at pre-defined terms. Reduces friction and captures expansion revenue without a sales cycle.
Capability adjustment provisions	Define how the commercial relationship evolves when the AI's capability materially improves. These provisions are novel and face the most resistance — but they are critical for closing the pricing lag force. Frame as: 'If we release a capability that materially improves the value you receive, we will offer it at pricing that reflects the improvement, and we will work with you to update our commercial relationship to reflect it.'
Anti-patterns to avoid	Hard commitment floors without expansion provisions (captures the commitment risk without the upside). Punitive overage pricing (creates adversarial dynamics when customers exceed their allocation). Long-term fixed pricing without any adjustment mechanism (guarantees the claw will open for the duration of the contract).

CASE STUDY: SNOWFLAKE
The Canonical Elastic Commercial Model at Scale

The structure	Snowflake charges per credit — a unit of compute and storage that scales with usage. As customers run more analytical workloads, more credits are consumed. As Snowflake's platform capability improves, the same workloads may require fewer credits, but the expanded scope of what customers do with Snowflake increases total consumption.
The elasticity in practice	Snowflake's NRR above 130% for multiple years is the commercial result of the elastic contract model: customers who sign an initial commitment expand their Snowflake usage as they deploy more analytical workloads, and the consumption-based pricing captures that expansion automatically. No renegotiation required. No sales cycle for expansion. The elastic contract grows with the customer's AI and analytics deployment.
The claw dynamics	Snowflake's per-credit pricing means that Snowflake improvements that make workloads more efficient (requiring fewer credits) could theoretically reduce per-workload revenue. Snowflake addresses this through the capability improvement dynamic: as Snowflake becomes more capable, customers use it for more workloads — total consumption grows even if efficiency per workload

	improves.
The design lesson	Snowflake's elastic model works because the credit definition is broad enough to capture all forms of usage growth: compute, storage, and data transfer all generate credit consumption. The elastic model must be designed so that growth in the customer's AI and data activity generates proportional revenue growth — not growth in one dimension while other dimensions are uncaptured.

CASE STUDY: DATADOG *Multi-Product Elastic Expansion*	
The elastic architecture	Datadog charges per host, per metric, and per log volume — with different rates for different product dimensions (infrastructure monitoring, APM, logs, security monitoring). As customers expand their Datadog deployment to more products, revenue grows proportionally with deployment breadth.
The NRR evidence	Datadog's NRR above 120% is the commercial result of the multi-product elastic expansion: customers who begin with infrastructure monitoring expand to APM, then to logs, then to security monitoring, then to network monitoring. Each expansion generates new revenue automatically from the pricing model — the contract allows expansion at catalog rates without requiring a new sales process.
The product launch elasticity	Datadog releases new products frequently (AI Observability, LLM Observability, Error Tracking). Each new product launch immediately expands the TAM from the existing customer base — customers who are already committed to Datadog's platform can add new products at catalog prices. The elastic contract architecture makes new product launches highly capital-efficient: the distribution exists, the relationship exists, the billing infrastructure exists. The new product is incremental revenue on the existing commercial relationship.
The claw dynamics	Datadog's per-host pricing means that host count growth drives revenue growth proportionally. As AI workloads add more hosts (more containers, more serverless functions, more AI inference infrastructure), Datadog's revenue grows with the AI economy's infrastructure expansion.

Chapter Nine — The Essentials

› Elastic contracts close the claw by capturing AI adoption growth and capability improvement automatically — without requiring a new sales cycle.

› Four mechanisms: consumption ratchets (volume discount + total revenue growth), expansion clauses (automatic terms for scope expansion), capability adjustment provisions (commercial framework for capability improvement value), and usage-based upsells (additional usage at catalog rates without new contracts).

› Snowflake's 130%+ NRR and Datadog's 120%+ NRR are the commercial results of elastic contract design at scale.

- › The critical design principle: the elastic contract must be designed so that growth in all dimensions of the customer's AI activity generates proportional revenue growth.
- › Capability adjustment provisions face the most customer resistance but address the most important claw force — make them a standard component of new enterprise contracts.

CHAPTER TEN

Closer 3 — Measurement First: Build the Instruments Before the Price

Why measurement infrastructure must precede pricing architecture. How to sequence the build.

Measurement first is the closing strategy that is most often neglected and most consistently critical to every other closing strategy. The commercial logic is simple: you cannot price what you cannot measure. You cannot close the claw with outcome anchoring if you do not have the infrastructure to verify outcomes. You cannot defend an elastic contract's capability adjustment provision without the data to demonstrate the capability improvement. You cannot benchmark your customer's AI performance without the data to establish the baseline and the metric.

Measurement must precede pricing architecture. This is the sequencing principle that distinguishes companies that successfully close their claw from those that continue to leave value on the table.

The sequencing failure is common: an AI company recognizes the open-claw problem, decides to implement outcome-based pricing, and attempts to implement it without the measurement infrastructure in place. The attempt fails because: the sales team cannot make the outcome commitment confidently (they do not have the evidence), the customer does not accept the outcome metric (it has not been validated), or the first billing dispute is unresolvable (the evidence chain does not exist). The company abandons the outcome pricing initiative and returns to subscription or per-token pricing — which perpetuates the claw.

The correct sequence:

Stage 1 (months 1–6): Instrument the product and customer integrations to capture the data required for value measurement. This is engineering work: adding measurement hooks to the AI system, building integrations with customer outcome systems, and establishing the data pipeline that brings outcome data alongside AI activity data.

Stage 2 (months 6–18): Build and validate the measurement methodology. Run the measurement on the first cohort of customers. Validate the attribution analysis. Establish confidence in the value estimates. Identify the customers whose deployments generate the clearest, most defensible evidence of value creation.

Stage 3 (months 18–24): Begin commercial model evolution with the evidence-ready customers. Introduce outcome-based pricing to the segment where the measurement is most confident. Use the first commercial cohort's results to build the evidence base for broader commercial model evolution.

Stage 4 (months 24+): Scale the outcome pricing model to the broader customer base, using the proven measurement methodology and the growing evidence base.

This is an 18–24 month investment in measurement before significant commercial return. Most AI companies are not willing to make this investment — they want revenue model improvements in the current quarter. The companies that are willing to make it are the ones that will have closed their claws while competitors are still measuring theirs.

The Gong.io measurement-first sequence is the clearest current example of a company that followed this discipline. Before Gong committed to outcome guarantees for enterprise customers, it spent 18 months building the analytical infrastructure that connected call conversation patterns to pipeline outcomes: win rate improvements, forecast accuracy improvements, rep ramp time reductions. The measurement came before the commercial commitment. When Gong introduced outcome-adjacent pricing with enterprise customers, the evidence base was solid enough to defend the conversation — and customers were receptive because they could see the data themselves in Gong's analytics dashboards.

The measurement infrastructure for specific AI product categories:

For productivity AI (coding assistants, document AI, knowledge management): time tracking integration that measures before/after time on specific task categories, combined with task quality metrics. The measurement challenge is establishing a clean baseline before deployment.

For process automation AI (workflow automation, data processing, order management): process timing metrics (cycle time before vs after AI deployment), error rate metrics, and cost-per-transaction metrics. These are typically already captured in the customer's operational systems — the measurement work is establishing the integration and the baseline.

For decision support AI (forecasting, analytics, recommendation systems): forecast accuracy metrics (comparing AI-assisted to human-only forecasts), decision quality metrics (were the recommended actions correct?), and outcome metrics (did the recommended decisions produce better outcomes?). These require careful attribution analysis to separate the AI's contribution from other factors.

For creative AI (content generation, design, marketing copy): quality scores (expert or customer assessment of AI-generated content), production time metrics, and downstream performance metrics (conversion rates for AI-generated marketing content vs human-generated). These require more subjective measurement and are harder to automate.

CLOSER 3: MEASUREMENT FIRST
The prerequisite. No commercial model evolution is sustainable without it. Build the instruments before you build the price.

The sequencing principle	You cannot price what you cannot measure. The measurement infrastructure must precede the pricing architecture — not follow it. The companies that have

	successfully closed their claws all built the measurement capability first and the commercial model second.
The four-stage sequence	Stage 1 (months 1–6): Instrument the product and build data pipelines. Stage 2 (months 6–18): Validate the measurement methodology on first customer cohort. Stage 3 (months 18–24): Begin commercial model evolution with evidence-ready customers. Stage 4 (months 24+): Scale to broader customer base with proven methodology.
The evidence threshold	Outcome pricing requires: (1) ≥ 50 customer deployments with validated measurement. (2) Attribution methodology validated with statistical confidence ≥ 85%. (3) At least 6 months of post-deployment measurement data. (4) Consistent results across diverse deployment contexts. Introducing outcome pricing before reaching this threshold risks the evidence being challenged and the commercial model being abandoned.
What happens without it	The common failure pattern: company recognizes open-claw problem → decides to implement outcome pricing → attempts implementation without measurement infrastructure → first billing dispute is unresolvable (evidence chain does not exist) → outcome pricing initiative is abandoned → company returns to subscription pricing. The measurement gap force continues unopened.
The ROI of measurement infrastructure	A measurement infrastructure investment of $500K–$2M in data engineering, integration, and statistical analysis typically enables commercial model changes that generate $5–20M in additional annual revenue. The payback period is 6–18 months from commercial model launch. The measurement infrastructure is the highest-return commercial investment most AI companies have not yet made.

CASE STUDY: GONG.IO
The Measurement-First Commercial Evolution

The discipline	Before Gong committed to outcome guarantees for enterprise customers, it spent 18 months building the analytical infrastructure connecting call conversation patterns to pipeline outcomes. The measurement came before the commercial commitment.
The measurement build	Gong's measurement infrastructure: (1) call analytics pipeline connecting conversation patterns (talk-to-listen ratio, question frequency, topic coverage) to deal outcomes (win/loss, cycle time, deal size); (2) rep performance attribution connecting coaching recommendations to specific behavior changes; (3) forecast accuracy measurement connecting Gong-assisted forecasts to actual quarterly results.
The evidence base before the pitch	When Gong's enterprise sales team began the outcome-adjacent commercial conversation, they arrived with: 200+ customer case studies with specific win rate improvements, a validated statistical model connecting Gong adoption

	patterns to pipeline conversion rates, and a customized value calculation for each prospect based on their team size, average deal value, and current win rate.
The commercial result	Gong's ability to quantify value creation enabled enterprise contracts of \$200K–\$500K annually for teams of 50–100 reps — contracts that were justified by documented ROI of \$1–3M in additional pipeline conversion from comparable deployments. The measurement infrastructure converted an abstract productivity claim into a specific, defended economic case.
The current state	Gong has measurement infrastructure that its competitors — Chorus.ai, Clari, Outreach, and others — cannot match. The measurement data is itself a competitive moat: Gong can make commercial commitments that competitors cannot make because Gong has the evidence base to defend them.

Chapter Ten — The Essentials

- › Measurement must precede pricing — this is the sequencing principle that distinguishes companies that close their claws from those that continue to leave value on the table.
- › The four-stage sequence: instrument (months 1–6), validate (months 6–18), begin evolution (months 18–24), scale (months 24+).
- › The evidence threshold before outcome pricing: ≥ 50 validated deployments, 85%+ attribution confidence, 6 months of post-deployment data, consistent results across contexts.
- › The ROI: \$500K–\$2M measurement investment → \$5–20M additional annual revenue → 6–18 month payback from commercial model launch.
- › Gong's measurement-first discipline is the case study: 18 months of analytics infrastructure building before the outcome-adjacent commercial conversation.

CHAPTER ELEVEN

Closer 4 — Moat Building: Defend Pricing Power Before the Claw Reopens

Data moats, workflow integration, trust infrastructure. Keeping the claw closed.

Moat building — defending pricing power before the claw reopens — is the long-term closer that ensures the commercial gains from Closers 1, 2, and 3 are sustained rather than competed away as the market matures.

The claw can be closed by implementing outcome pricing, elastic contracts, and measurement infrastructure — and then reopened by competitive dynamics that compress the pricing back toward commodity levels. If the vendor's AI capabilities are broadly replicable (as Book 6 analyzed), competitors will enter with comparable capabilities at lower prices, and the value-based pricing that justified the closed claw will be pressured downward.

Moat building for claw sustainability requires investing in the specific competitive advantages that protect pricing power after the claw is initially closed. Three moat types are most relevant:

Data moats that deepen with claw closure. When outcome pricing generates outcome data, that data — the record of what the AI accomplished, under what conditions, with what results — is itself a valuable asset. The vendor that has measured 10 million contract reviews has a dataset that is extraordinarily valuable for training better contract review AI. The better the AI, the more valuable the reviews become, the more defensible the price. The data generated by closed-claw commercial operations becomes the training asset for the next capability generation — a compound loop that favors the company that closes its claw first.

Workflow integration that locks closed-claw value. When outcome pricing requires deep integration with customer systems (to verify outcomes), that integration is itself a switching cost. The customer's system of record for contract review outcomes is now integrated with Harvey's billing system. Migrating to an alternative contract review AI requires re-integrating the verification system, re-establishing the baseline, and accepting a period of measurement uncertainty. The closed claw's verification infrastructure becomes the workflow lock that defends the pricing.

Customer evidence accumulation that creates the next commercial opportunity. Each successfully closed claw — each customer where outcome pricing is working, where the measurement is validated, and where the customer is satisfied with the value share — is a reference case for the next customer. The company with 200 documented case studies of AI-measured ROI is in a fundamentally different commercial position than the company with 20. The closed claw's evidence base compounds into the next commercial generation.

Specific moat-building investments that protect claw closure:

Anthropic's investment in constitutional AI and safety documentation is a trust moat that protects its enterprise pricing. Enterprise customers who need AI deployed in sensitive contexts are willing to pay a premium for AI systems with documented safety practices and responsible AI governance. This trust moat is not directly related to AI capability — it is an accountability and governance moat that defends pricing against lower-cost alternatives that cannot offer comparable governance documentation.

Salesforce's customer data integration in its AI products creates a data moat that compounds with usage. Every interaction managed through Salesforce's AI generates data about sales patterns, customer behavior, and commercial outcomes that is used to improve Einstein AI features. The customers who have used Salesforce AI the longest have the most data informing their AI recommendations — which makes the AI more valuable for them and makes switching to an alternative more costly.

ServiceNow's workflow certification and compliance infrastructure is an accountability moat that protects its premium pricing in enterprise workflow AI. ServiceNow's AI features are deployed in the workflows of 85% of the Fortune 500. The compliance certifications, the audit trails, and the enterprise governance infrastructure that ServiceNow has built make its AI-assisted workflows the lowest-risk option for enterprise operations — a trust moat that is not easily replicated by lower-cost AI alternatives.

CLOSER 4: MOAT BUILDING

Close the claw, then defend it. Competitive pressure will try to reopen the gap. Build the moats that make the closed position sustainable.

The sustainability challenge	Closing the claw through outcome pricing, elastic contracts, and measurement infrastructure creates a commercial position that generates significantly more revenue per customer. This commercial success attracts competitors who will attempt to compete at lower prices. The moat-building closer ensures the closed

	position is defensible against this competitive pressure.
Three moat types for claw sustainability	(1) Data moats: the outcome measurement data generated by closed-claw commercial operations becomes the training asset for the next capability generation. (2) Workflow integration: the verification infrastructure required for outcome billing creates switching costs that defend the commercial position. (3) Customer evidence accumulation: each successfully closed claw is a reference case that compounds into the next commercial generation.
The data moat loop	Close claw → generate outcome data → train better AI on outcome data → AI improves → outcomes more valuable → pricing power increases → close claw further. This loop is the commercial expression of the AI investment flywheel. The company that starts the loop first compounds fastest.
The workflow integration moat	Outcome billing requires integration with the customer's system of record. This integration is itself a switching cost: to switch vendors, the customer must re-integrate the verification system, re-establish the baseline, and accept measurement uncertainty during the transition. The closed claw's operational infrastructure becomes the workflow lock.
The trust moat	Companies that have closed their claws through outcome billing have demonstrated commercial accountability — they have been willing to put their revenue at risk on their AI's performance. This demonstrated accountability is a trust signal that attracts customers who value commercial credibility and that defends the pricing premium against less accountable alternatives.

CASE STUDY: ANTHROPIC
The Safety and Trust Moat in Foundation Model Pricing

The commercial context	Anthropic operates in the foundation model market — a market facing significant commodity pressure as multiple capable models compete on price. Per-token prices have been declining across the industry. The commodity pressure would normally force Anthropic into a race-to-zero on pricing.
The trust moat	Anthropic's investment in Constitutional AI, its Responsible Scaling Policy, and its extensive safety research documentation creates a trust moat that is specific to the AI era: enterprise buyers deploying AI in sensitive contexts (healthcare, legal, financial services, government) are willing to pay a premium for AI systems with documented safety practices, transparent model cards, and responsible AI governance frameworks.
The commercial evidence	Anthropic's enterprise contracts (Claude API Enterprise) are priced above commodity API access for comparable capability — a premium that is sustained by the trust moat. Enterprise customers whose IT security and compliance teams require AI governance documentation are not in the commodity market: they are in the accountability market, where Anthropic's documented safety practices create a commercial moat.

The claw sustainability	Anthropic's safety investment is not just a trust signal — it is infrastructure that makes it harder for competitors to match the commercial position. Building comparable safety documentation, Constitutional AI training methodology, and responsible AI governance takes years. The moat is real and currently growing, as AI safety concerns increase enterprise attention to AI governance.

CASE STUDY: SALESFORCE
The CRM Data Moat and the Einstein Flywheel

The moat architecture	Salesforce's CRM data — 150,000+ customers' commercial interaction histories, deal patterns, win/loss records, and customer behavior data — is the training asset for Einstein AI features. The moat compounds: more customers → more data → better Einstein AI → better outcomes → more customers.
The claw sustainability implication	When Salesforce closes its commercial claw on AI features (through Agentforce pricing at $2/conversation, through Einstein outcome analytics, through Copilot capabilities priced at premium tiers), the revenue generated funds the AI development that deepens the data moat. Competitors who are at lower capture rates generate less revenue to invest in AI — which allows the data moat to grow while their capability lags.
The workflow integration moat	Salesforce's CRM integrations — with email (Gmail, Outlook), calendar, marketing automation, and customer service platforms — create workflow switching costs that are independent of AI feature pricing. The data flows that these integrations create are also the training data for Einstein AI. The workflow moat and the data moat are the same infrastructure, viewed from two perspectives.
The commercial result	Salesforce's 100%+ NRR across multiple years and 75%+ gross margins are the commercial results of closed-claw commercial operations supported by compounding data and workflow moats.

Chapter Eleven — The Essentials

› Moat building is the sustainability closer: without it, competitive pressure reopens the claw that Closers 1–3 have closed.

› Three moat types: data moats (outcome data trains better AI), workflow integration (billing infrastructure creates switching costs), customer evidence (reference cases compound into commercial position).

› The data moat loop: close claw → generate outcome data → train better AI → AI improves → outcomes more valuable → pricing power increases → close claw further.

› Anthropic's safety trust moat demonstrates that non-capability moats can sustain premium pricing in commodity capability markets.

› Salesforce's Einstein data flywheel is the compounding data moat at maximum scale: 150,000 customers' commercial data continuously improving the AI that justifies premium pricing.

PART FOUR

The Claw and the Future

Agent-to-agent commerce. Regulated industries. The manifesto for closing.

CHAPTER TWELVE

The Open Claw in Agent-to-Agent Commerce

When AI buys from AI, the claw dynamic changes completely. The economics of machine procurement.

Agent-to-agent commerce — the emerging category where AI agents purchase services from other AI agents without human involvement in individual transactions — creates a fundamentally different open-claw dynamic than human-mediated commerce.

The conventional open-claw problem is a mismatch between AI advancement speed and human commercial process speed. Humans sign annual and multi-year contracts. Procurement processes move at human decision-making pace. Commercial model evolution requires human negotiations. All of this creates the structural lag between capability and capture.

When AI agents are the buyers, the commercial dynamics change completely:

AI agents transact at machine speed. A human enterprise signs a three-year contract. An AI agent can evaluate, negotiate, and execute a commercial agreement for a specific service in milliseconds — selecting the vendor with the best price-performance ratio for the specific task at hand, every time a task is required.

AI agents have perfect price memory and perfect comparison capability. A human buyer evaluates vendors periodically and accepts the friction of switching. An AI buyer evaluates every vendor for every transaction, executes with the best price-performance option, and switches immediately if a competitor offers better value.

AI agents can negotiate continuously. A human signs a contract and locks in pricing for the term. An AI agent can continuously renegotiate based on current market conditions, competitive alternatives, and the specific requirements of the current task.

These dynamics suggest that the open-claw problem in agent-to-agent commerce takes a different form. The traditional open claw is about capturing value that currently exists in commercial relationships with humans who accept pricing that is below the value created. In agent-to-agent commerce, the buyers will not accept pricing below the value received — they will instantly switch to alternatives. The claw in A2A commerce opens in the opposite direction: the pressure is on the vendor to continuously demonstrate value equal to or greater than the price, or lose the agent buyer immediately.

This creates a different commercial imperative for AI vendors serving agent buyers: the commercial model must be outcome-based from the start (agents will not pay for access to capability; they will pay for results), the pricing must be continuously competitive (agents will continuously compare alternatives), and the value measurement must be real-time (agents will base their purchasing decisions on measured, verified performance metrics, not on vendor claims or historical case studies).

Specific examples of early agent-to-agent commerce patterns:

AI-powered research agents that purchase web search, data extraction, and analysis services from specialized AI providers represent the earliest form of agent-to-agent commerce. The orchestrating agent selects and pays for these services based on real-time price-performance comparisons. The vendors in this ecosystem (Browserbase, Exa, Firecrawl) are already experiencing the A2A commerce pricing dynamics: agents compare them continuously, switch based on performance, and the market prices are determined by competitive dynamics rather than by annual contract negotiations.

API marketplaces (RapidAPI, Apilayer) that serve AI agents as buyers are beginning to see AI agent usage that differs from human API consumption patterns: higher volumes, more consistent consumption, immediate switching on performance degradation, and price sensitivity calibrated to the task value rather than to budget cycles. The vendors on these marketplaces are being forced to price based on continuous competitive dynamics rather than on the procurement cycles of their historical human buyers.

The MCP server ecosystem — the AI agent service marketplace enabled by the Model Context Protocol described in Book 4 — is designed from the ground up for agent-to-agent commerce. MCP servers that provide specialized capabilities (search, code execution, data analysis, API integration) are accessed by AI agents at the moment they need those capabilities, priced per use, and selected based on real-time performance characteristics. The commercial model for MCP servers is inherently outcome-based, inherently elastic, and inherently competitive.

The commercial infrastructure for A2A commerce requires specific design choices that differ from human-mediated commercial infrastructure: real-time price publication (agents need to query current prices for specific task types, not annual catalog prices), performance attestation (agents need verified performance metrics to make purchasing decisions), micropayment settlement (A2A commerce often involves very small transaction values that invoice-based settlement cannot handle), and automated dispute resolution (the volume of A2A transactions makes human dispute review impractical).

A2A Commerce vs Human-Mediated Commerce — Commercial Dynamics

Dimension	Human-mediated commerce	Agent-to-agent commerce	Implication for claw
Transaction speed	Annual/multi-year contracts; quarterly	Millisecond selection, execution,	Pricing lag force disappears —

	review cycles	and settlement per transaction	agents re-price at every transaction
Price memory	Buyers accept historical pricing anchors; switching friction creates price stickiness	Agents have perfect price memory and zero switching friction; buy at best current price every time	Race-to-zero force intensifies — no anchoring effect from historical pricing
Value measurement	Buyers often cannot measure AI value precisely; value-based pricing requires vendor evidence	Agents can measure and compare performance in real time across vendors	Measurement gap force disappears — agents measure continuously
Contract lock	Multi-year contracts prevent commercial model adjustment	No long-term contracts in machine-speed commerce; every transaction is a fresh evaluation	Contract lock force disappears — replaced by continuous competitive evaluation
Competitive pressure	Limited by buyer switching costs and procurement friction	Zero switching costs; instant competitive comparison; instant substitution	Commercial model must be continuously competitive or lose transactions immediately

THE A2A PREPARATION IMPERATIVE

Build A2A-ready commercial infrastructure before the transition forces it — the commercial infrastructure for human-mediated contracts is not ready for machine-

speed transactions

The infrastructure required for A2A commerce readiness: (1) Real-time pricing APIs that publish current prices for specific task types — not annual catalog prices but real-time prices that reflect current competitive positioning. (2) Performance attestation APIs that publish verified performance metrics — agents need objective performance data, not vendor claims. (3) Micropayment settlement infrastructure that handles high-frequency, low-value transactions without the overhead of invoice-based billing. (4) Automated dispute resolution that scales to machine-speed transaction volume. The time to build this infrastructure is now — before A2A customers require it. The infrastructure is 18–24 months to build. A2A commerce is already beginning in some markets.

Chapter Twelve — The Essentials

- A2A commerce inverts the claw dynamics: human-mediated commerce creates pricing lag; machine-mediated commerce eliminates it.
- In A2A commerce, the four claw-opening forces disappear — replaced by continuous competitive evaluation that makes pricing pressure immediate rather than lagged.
- The A2A claw opens in the opposite direction: vendors who deliver less value than they price lose transactions instantly; vendors who deliver more value have room to close the gap through better pricing.
- The commercial infrastructure for A2A commerce: real-time pricing APIs, performance attestation, micropayment settlement, automated dispute resolution.
- Build A2A-ready commercial infrastructure now — 18–24 months before A2A customers require it, because the infrastructure must precede the commercial model, not follow it.

CHAPTER THIRTEEN

The Claw in Regulated Industries: Healthcare, Finance, Legal

How regulatory constraints widen the claw — and how to close it within compliance.

The open-claw effect in regulated industries — healthcare, financial services, legal — is more severe than in unregulated technology markets because regulatory constraints widen the claw while simultaneously making it harder to close.

The widening mechanism: regulatory compliance requirements slow the deployment of new AI capabilities, creating a longer lag between AI capability availability and commercial deployment. A healthcare AI that can diagnose conditions with expert-level accuracy cannot be deployed commercially until it has completed the FDA clearance process — which can take 18–36 months. During that period, the AI capability is improving in the research environment, but the commercial deployment is frozen at the baseline capability level. The claw opens during the regulatory review period and must be closed after the review is complete.

The closing difficulty: the same regulatory frameworks that slow deployment also impose specific constraints on outcome-based pricing that are difficult to navigate. In healthcare, for example, outcomes are subject to complex regulatory frameworks governing diagnostic claims: an AI that says "this imaging scan shows a 94% probability of malignancy" is making a diagnostic claim that is subject to FDA oversight as a medical device. Charging per accurate diagnosis is potentially charging for a regulated medical service — with liability implications that make straightforward outcome pricing commercially complicated.

However, regulated industries are also the contexts where the value of AI capabilities is highest and where the commercial opportunity from closing the claw is greatest.

Healthcare AI open-claw examples:

Nuance (acquired by Microsoft) in clinical documentation AI represents a company that partially closed its claw through the measurement-first approach. Nuance's Dragon Ambient eXperience (DAX) product reduces physician documentation time by an average of 50% — from approximately 4 hours per day to 2 hours per day, freeing 2 hours of physician time that is currently the most valuable and most constrained clinical resource in healthcare. The value created per physician per day is approximately $500–800 (2 hours × $250–400 per hour physician time).

Nuance's pricing for DAX has moved progressively from per-provider subscription (which did not capture the documentation time value) toward capacity-based models that better reflect the value created. The claw closure challenge in healthcare is the regulatory and procurement complexity: hospital procurement of physician-facing AI tools involves complex multi-stakeholder evaluation processes that slow commercial model evolution.

The measurement-first solution in healthcare: before Nuance could implement value-based pricing, it needed to measure what was actually happening to physician documentation time after DAX deployment. The measurement data — average reduction in documentation time by specialty and practice setting — is now robust enough to support value-based pricing conversations with health system CFOs. The 18-month measurement investment has created the evidence base for claw closure.

Financial services AI open-claw examples:

Addepar, the wealth management analytics platform, operates in a regulatory environment where the output of its AI — portfolio analytics, performance attribution, client reporting — must comply with SEC rules on performance reporting and disclosure. The AI's capability has advanced significantly since Addepar's initial deployment, but the pricing model (primarily AUM-based or per-account) does not automatically capture the capability improvement.

The claw closure path for Addepar is outcome-adjacent pricing: connecting the price to specific value metrics that the regulatory framework does not constrain. Better risk-adjusted portfolio performance attribution, faster client reporting generation, and improved compliance monitoring are value metrics that can be priced without creating regulatory complexity. The outcome pricing focus on efficiency and quality improvement, rather than on investment performance, allows outcome-based pricing to coexist with the investment performance regulatory constraints.

Legal AI open-claw examples:

Casetext (acquired by Thomson Reuters) in legal research AI deployed an AI that dramatically accelerated legal research — reducing the time for complex research tasks from hours to minutes. The value creation was clear: law firms and corporate legal departments pay associates $50–200/hour for research work, and an AI that completes comparable research in 10% of the time is creating obvious value.

The claw closure challenge in legal AI is the unauthorized practice of law constraint: in most US jurisdictions, providing legal advice for compensation requires a law license. An AI that provides specific legal research conclusions is approaching the boundary of providing legal advice. Outcome-based pricing for legal research AI must be structured around the efficiency value (time saved on research) rather than the quality value (accuracy of legal conclusions) to avoid crossing the unauthorized practice line.

Thomson Reuters' acquisition of Casetext for $650M reflects the strategic value of closing the claw in legal AI: Casetext's measurement infrastructure for legal research efficiency and its established outcome-adjacent pricing model represented a commercial position that justified a substantial acquisition premium over the standalone AI capability value.

Regulated Industry Claw — Specific Constraints and Closers

Industry	Regulatory constraint	How it widens the claw	Closing approach
Healthcare	FDA clearance required for AI diagnostic claims (510(k), De Novo, PMA pathway depending on risk level). 18–36 month process.	AI capability improves during FDA review period; commercial deployment frozen at pre-clearance capability level; capability-capture gap widens during review	Price for efficiency outcomes (physician time saved, documentation quality improved) rather than diagnostic performance — efficiency outcomes are not diagnostic claims and do not require FDA clearance
Financial services	SEC regulation of performance claims for investment AI; FINRA rules on AI advisor	Cannot price outcomes tied to investment performance (regulatory	Price for operational efficiency (analyst time reduction, report generation speed, compliance documentation accuracy) rather than investment performance

	recommendations; bank regulatory approval for AI in credit decisions	prohibition); must price efficiency or compliance outcomes	outcomes
Legal	Unauthorized practice of law regulations in most US jurisdictions; similar restrictions globally	Cannot price for the quality of legal advice delivered; legal conclusions are regulated services	Price for research efficiency (time to relevant precedent, case law comprehensiveness) and document processing (review time, issue identification rate) rather than for the legal advice quality the research supports
Healthcare (clinical)	HIPAA data handling requirements; CMS documentation standards; state medical board regulations	Data governance requirements increase implementation complexity and slow commercial deployment; CMS documentation standards constrain output format flexibility	Build CMS-compliant documentation workflow into the commercial offering; HIPAA compliance certification enables pricing premium; CMS compliance certification enables healthcare system procurement

Chapter Thirteen — The Essentials

› Regulated industries have wider open claws because regulatory review periods freeze commercial deployment while AI capability continues to improve.

› Healthcare: price for efficiency outcomes (physician time), not diagnostic performance — efficiency is outside FDA's medical device framework.

› Financial services: price for operational efficiency and compliance documentation, not investment performance — performance claims are regulated.

› Legal: price for research efficiency and document processing, not legal advice quality — advice quality implicates unauthorized practice regulations.

› Compliance certification (HIPAA, SOC 2, CMS) is itself a closing mechanism: it enables procurement in sectors that require certification and commands a trust premium.

CHAPTER FOURTEEN

The Manifesto for Closing the Claw

The eight commitments every AI company should make. The closing argument.

The eight commitments of claw closing are the operating principles that distinguish organizations serious about capturing the value their AI creates from those that are content to leave it on the table.

THE EIGHT COMMITMENTS OF CLAW CLOSING	
1	Measure before you price. No commercial model for AI capability improvement is defensible without measurement infrastructure. Invest in measurement before commercial model evolution. Use the measurement data to build the value anchor before attempting to capture it.
2	Price outcomes, not inputs. Per-token, per-seat, and per-API-call pricing are all input pricing. They measure what went in, not what came out. Price what the AI accomplishes. When you cannot price outcomes directly, price the output most closely correlated with the outcome you create.
3	Build elasticity into every contract. No contract should be signed without expansion provisions. No AI deployment should grow without generating proportional revenue. Design every contract for the AI's success — make the commercial terms reflect the upside of adoption rather than only the risk of the initial commitment.
4	Communicate AI improvement to customers. When your AI gets materially better, tell your customers. Show them the evidence. Use the improvement announcement as the context for a commercial conversation about whether the current pricing reflects the current value. The improvement communication is the commercial catalyst.
5	Close the claw in your most defensible segment first. Find the customers where the measurement is clearest, the value is most defensible, and the relationship is strongest — and close the claw there first. Use those customers as the evidence base for broader claw closure.
6	Build the data moat while closing the claw. The data generated by outcome-measurement infrastructure is a strategic asset. Store it, analyze it, and use it to train better AI models. The measurement infrastructure and the data moat are the same investment.
7	Move to A2A-ready commercial infrastructure before the transition forces it. Agent-to-agent commerce is coming. The commercial infrastructure for human-mediated contracts is not ready for machine-speed transactions. Build real-time pricing APIs, performance attestation, and micropayment settlement now.
8	Report the claw to the board quarterly. The gap between AI value created and AI value captured is a governance metric. Make it visible. Boards that do not see the claw cannot

govern for its closure. The Claw Dashboard belongs in the board pack.

EXTENDED CASE STUDIES

Three Views of the Open Claw

OpenAI · Microsoft Copilot · Anthropic — three different positions in the capability-capture dynamic.

CASE STUDY A

OpenAI: The Claw Opening in Real Time

Five years of capability improvement against a per-token pricing model that systematically underprices the value created.

OpenAI's pricing history is the clearest large-scale example of claw opening and partial closure in the AI economy — a history that tracks the capability-capture gap over five years of unprecedented AI advancement.

In 2020, when OpenAI launched GPT-3 for commercial access, the initial pricing was approximately $0.06 per 1,000 tokens. This price was set against the capability of GPT-3 — a model that could generate coherent text, complete prompts, and perform simple question-answering tasks. The commercial model was per-token: pay for what you consume.

Between 2020 and 2024, the open claw began its opening:

GPT-3.5 (2022) was 10× better than GPT-3 at coding and instruction following, but the per-token price was comparable. Customers who were using GPT-3 for tasks that GPT-3.5 could handle dramatically better were getting dramatically more value at the same price.

GPT-4 (2023) was qualitatively better at complex reasoning, analysis, and professional work — capable of passing bar exams, medical licensing exams, and coding interviews. GPT-4's pricing was higher ($0.03–0.06 per 1K tokens for 8K context), but the value improvement was not linearly reflected in the price. A customer using GPT-4 for complex contract analysis was getting value comparable to a junior attorney's work at a price comparable to a cheap automation tool.

GPT-4 Turbo (late 2023) added 128K context window at lower cost than the original GPT-4. The value of being able to analyze an entire 100-page legal agreement in a single prompt was dramatically higher than the value of analyzing 8 pages — but the price per token actually declined.

GPT-4o (2024) was faster, cheaper, and better than GPT-4 across most benchmarks. OpenAI lowered the price — $5 per million input tokens vs $30 for GPT-4. The value per token increased substantially; the price per token decreased.

This pricing history is a textbook claw-opening trajectory. Each capability generation delivered more value; the price structure systematically underpriced relative to value as capabilities improved. OpenAI's revenue grew substantially as volume increased, but the capture rate (the percentage of value created that was captured as revenue) almost certainly declined over this period as capability improvements outpaced pricing adjustments.

The partial claw-closure attempt came with ChatGPT Plus ($20/month), the Teams ($25/user/month) and Enterprise (custom) offerings, and the GPT Store monetization. Each tier captured more value from specific customer segments by matching pricing to the value those customers received. But the core API pricing continued to decline per token even as per-token value increased.

OpenAI's commercial evolution toward outcome-correlated products (GPT-4 API with structured outputs, real-time voice, code interpreter) represents the direction of claw closure at the foundation model layer. These capabilities are priced at premiums to base token pricing that better reflect the value they create — code interpreter charges above base token rates because it is a tool that accomplishes more per token, not just a model that generates text per token.

The commercial lesson from OpenAI's trajectory: per-token pricing is structurally claw-opening in a world of AI efficiency improvement. Every capability generation that reduces cost-per-equivalent-output while increasing value-per-token widens the gap between what is charged and what is created. The commercial model must evolve from input pricing (per token, per API call) toward output pricing (per task accomplished, per outcome verified) to close the claw structurally.

CASE STUDY B

Microsoft Copilot: Systematic Tier-Based Claw Closure

From $10/user/month GitHub Copilot to a multi-tier enterprise AI ecosystem — the most detailed current example of planned claw closure.

Microsoft's Copilot commercial trajectory — from GitHub Copilot's $10/user/month launch in 2022 to the multi-tier enterprise AI ecosystem of 2024-2025 — is the most detailed current example of a major software company systematically working to close its open claw across multiple product dimensions.

The original claw: GitHub Copilot launched in June 2022 at $10/user/month for individuals (with a $100/year annual option), $19/user/month for businesses. The capability at launch: AI-powered code completion that improved developer productivity by an estimated 30–50% in studies. The value created: at a developer fully-loaded cost of $200/day, 30% productivity improvement is $60/day in time value — $1,200/month per developer. The claw at launch: $10/month charged versus $1,200/month in documented time value = 0.8% capture rate. A significant open claw.

The systematic closure program:

Step 1 — Volume (2022–2023): GitHub grew Copilot to 1M+ paid subscribers at the initial pricing. This was capturing 0.8% of developer productivity value, but at scale. The commercial priority was adoption — getting into developer workflows before competitors.

Step 2 — Tier expansion (2023): GitHub introduced Copilot Business at $19/user/month, differentiated from Individual by organizational management features and data privacy controls (code not used for model training). The $9 premium over Individual was modest — it was not a value-based price increase but a feature-differentiation increase. The claw remained wide.

Step 3 — Enterprise tier (2024): GitHub Copilot Enterprise at $39/user/month ($480/year) was the first pricing tier that approached meaningful value capture. The Enterprise tier adds: organization-specific fine-tuning (Copilot understands the organization's codebase and conventions), deeper IDE integration, GitHub.com integration for code search and pull request assistance, and audit logs. At $39/user/month against $1,200/month in value, the Enterprise tier captures 3.25% — still a wide claw, but meaningfully closer than the Individual tier.

Step 4 — Adjacent commercial expansion (2024-2025): Microsoft 365 Copilot ($30/user/month for enterprise customers) extended AI assistance to Word, Excel, PowerPoint, Outlook, and Teams. The value proposition: each knowledge worker saves 1–2 hours/day through AI assistance in their primary productivity tools. At a $100K/year knowledge worker cost, 1 hour/day savings = $12,500/year = $1,042/month in value. Microsoft 365 Copilot charges $30/month = 2.9% capture rate. Still a wide claw, but the commercial model is moving toward the value.

Step 5 — The ongoing closure program: Microsoft's AI commercial roadmap includes Copilot Studio (allowing organizations to build custom AI agents on the Copilot framework), Copilot in Security (AI-assisted security operations at higher price points), and Copilot for specialized roles (Sales Copilot, Service Copilot, Finance Copilot) priced at the value of the specific role's productivity improvement rather than at the general Copilot price.

The commercial insight from Microsoft's trajectory: the claw closure program works through tier stratification — successive tiers that capture progressively more value from segments where the value is greatest and most defensible. Each tier closes the claw for one segment without requiring a global price increase that would face resistance from the broader customer base.

CASE STUDY C

Anthropic: Foundation Model Claw and Enterprise Layer Opportunity

The structural open claw at the foundation model layer — and the enterprise pricing opportunity that addresses it.

Anthropic's commercial evolution provides a distinctive perspective on the open-claw problem because Anthropic operates simultaneously as a foundation model provider (with the classic per-token pricing claw problem) and as an enterprise AI provider (with the opportunity to price toward business outcomes rather than toward API calls).

The foundation model layer claw: Claude's API pricing has followed the industry pattern of capability improvement without proportional price increases. Claude 3 Haiku at $0.25 per million input tokens is dramatically more capable than the first Claude APIs at comparable or higher prices. The per-token price has declined while the per-token value has increased. The foundation model layer claw is wide and structurally difficult to close because per-token pricing is the industry standard and customers expect prices to decline as models improve.

The enterprise layer opportunity: Anthropic's enterprise Claude offering (Claude.ai Enterprise) is where the claw closure opportunity is most significant. Enterprise customers using Claude for complex analytical tasks — processing large document sets, conducting multi-step research, generating nuanced professional communications — are deriving value that far exceeds the per-token cost. The enterprise pricing structure (custom pricing based on usage volume and enterprise requirements) allows Anthropic to price closer to the value created for specific high-value use cases.

The Claude Code commercial frontier: Claude Code, Anthropic's AI coding agent, represents Anthropic's most explicit move toward outcome-adjacent pricing. Claude Code is priced per interaction rather than purely per token — the pricing acknowledges that the value of a completed coding task is not proportional to the tokens consumed to complete it. This is a partial step toward outcome anchoring: pricing that is closer to the output (a completed task) than to the input (tokens processed).

The measurement challenge at the foundation model layer: Anthropic's measurement problem is that the API's customers build products on top of Claude, and the value those products create is not directly visible to Anthropic. A company that uses Claude API to power a contract review service creates substantial value for end customers — but Anthropic only sees the API consumption, not the contract reviews completed or the legal disputes avoided. The intermediate layer between API and end value makes measurement difficult and outcome pricing at the API layer impractical.

The commercial evolution direction: Anthropic's commercial trajectory suggests a move toward specialized API products for specific high-value use cases — APIs that are priced based on the specific capability they provide (structured output generation, long-document analysis, multi-step reasoning) rather than purely on token consumption. This is a partial claw closure strategy: specialized capabilities are priced at the value of what they accomplish rather than purely at the cost of what they consume.

CLOSING

The Claw Opens Silently

Every quarter you wait, it widens. Build the commercial infrastructure to close it.

Framework F22 — The Claw Dynamics Model + The Four Closers

Framework F22 — The Claw Dynamics Model + The Four Closers — is the definitive commercial framework for the AI capability-capture gap.

The Claw Dynamics Model has three components:

The capability curve: AI capability growing exponentially — each capability generation significantly more powerful than the previous, with the trajectory accelerating as models train on more data, as architectures improve, and as inference costs fall.

The capture curve: Revenue capture growing arithmetically — constrained by annual contract cycles, by measurement infrastructure that lags capability deployment, by procurement processes that operate at human speed, and by the legacy pricing structures inherited from the AI's less capable past.

The gap: The continuously widening space between the two curves — the commercial value created by AI capability that is not being captured in revenue. This is the open claw. It opens silently, every quarter, in every AI company that has not built the commercial infrastructure to close it.

The Four Closers are the specific commercial strategies that close the claw and keep it closed:

Closer 1 — Outcome Anchoring: Price what the AI accomplishes. Connect the commercial model to verified, economically valued outcomes. When the AI improves, the value of each outcome may increase, but the price structure remains appropriate because it is anchored to the outcome, not to the AI capability that produces it.

Closer 2 — Elastic Contracts: Build expansion into the commercial infrastructure. Consumption ratchets, expansion clauses, capability adjustment provisions, and usage-based upsells ensure that AI adoption growth and capability improvement generate revenue automatically rather than requiring renegotiation.

Closer 3 — Measurement First: Build the instruments before the price. The value measurement infrastructure is the prerequisite for every other claw-closing strategy.

Measure before you commit. Evidence before you anchor. Verification before you invoice.

Closer 4 — Moat Building: Defend pricing power before the claw reopens. Data moats, workflow integration, customer evidence accumulation, and trust infrastructure ensure that the commercial gains from claw closure are sustainable against competitive pressure.

Framework F22 — The Four Closers Reference

Closer	What it addresses	How it closes the claw	Primary force it closes	Success indicator
1. Outcome Anchoring	Per-token and per-seat pricing that fails to capture value as AI capability improves	Prices the output (outcome) rather than the input; capability improvement increases outcome reliability and volume rather than eroding revenue	Pricing lag force: per-input pricing makes capability improvement a customer benefit, not a revenue growth driver	Revenue per customer grows as AI capability grows; NRR reflects outcome delivery expansion
2. Elastic Contracts	Fixed contracts that lock pricing at a point in time while AI capability and customer deployment grow	Consumption ratchets, expansion clauses, and capability adjustment provisions ensure revenue grows with AI adoption and improvement	Contract lock force: long-term contracts lock pricing at historical capability levels	NRR consistently above 115%; revenue grows from existing customers without new sales cycles
3. Measurement First	Absence of value measurement infrastructure that prevents	Builds the evidence chain from AI activity to business	Measurement gap force: invisible value cannot be priced	Outcome pricing conversations succeed; commercial

	outcome anchoring and value-based pricing	outcome to dollar value; creates the foundation for all other closers		commitments are defended; billing disputes are resolved by evidence
4. Moat Building	Competitive pressure that reopens the claw after it has been commercially closed	Data moats from outcome data, workflow integration moats from billing infrastructure, trust moats from accountability evidence all defend the closed position	Race-to-zero force: competition drives pricing toward commodity levels	Pricing premium sustained against commodity competition; customer switching rate below market average

The open claw opens silently.

There is no moment when the claw clearly announces itself. The AI improves incrementally, in ways that are visible to the engineering team but not to the contracts team. The value delivered grows quarter over quarter, in ways that are visible in customer outcomes but not in revenue reports. The gap between capability and capture widens in the background, invisible until someone asks the right question.

The right question is not "how much revenue are we generating?" It is "how much value are we creating, and what fraction of it are we capturing?" The ratio between those two numbers is the state of your claw. If the ratio is declining — if capture as a fraction of value is falling even as revenue grows in absolute terms — the claw is opening.

The companies that ask this question and act on the answer will build the most valuable AI businesses of the next decade. Not because they are building better AI — there are many companies building excellent AI. But because they are building better commercial infrastructure around their AI: the measurement systems, the outcome pricing models, the elastic contracts, and the competitive moats that allow them to capture a rising fraction of the rising value their AI creates.

The compounding math is relentless. A company with a 20% capture rate and a 15% monthly AI capability improvement compounds at 1.20 × 1.15 = 1.38 per month on value creation, while revenue compounds at 1.20 × 1.00 = 1.20 per month if the capture rate does not improve. After 12 months, the value creation has increased by 5.35×; the revenue has increased by 8.9×... but the gap is $5.35 – $8.9/5 = still $4.56 uncaptured per dollar of original value. The absolute revenue grows; the capture rate gap widens.

The company that closes its claw — that matches its revenue compounding to its value compounding — captures the full potential of its AI investment. The company that does not watches an ever-larger fraction of its AI value become someone else's benefit.

The open claw opens silently. Every quarter you wait, it widens. The question is not whether you have an open-claw problem. It is how big yours already is — and whether you are going to close it.

> ***"The open claw opens silently. Every quarter you wait, it widens. The question is not whether you have an open-claw problem. It is how big yours already is — and whether you are going to close it."***

The AI Economy Monetization Series continues in Book Nine:

Customer Success in the AI Economy