

A DIAGNOSTIC FRAMEWORK

The AGENT-FIT™ Framework

**Six lenses for identifying
high-value Agentic AI use cases.**

Finance

HR

Sales

+ worked examples

THE HARD TRUTH

**Most agentic AI
pilots don't fail in
~~execution.~~**

They fail in selection.

The hardest question in agentic AI isn't
"how do we build it?"

It's: **"should we build THIS one at all?"**

THE FRAMEWORK

Six Lenses. One scoreable conversation.



Score each **1–5** · Plot on Value × Feasibility · Sequence accordingly

LENSES 1 — 3

Is the use case **well-formed**?



LENS 01

Outcome Clarity

Can success be stated in one sentence?

- ▶ Measurable in \$, time, or risk
- ▶ Owned by a named accountable executive
- ▶ Baseline and target both quantified



LENS 02

Process Topology

Is the work agent-shaped, not script-shaped?

- ▶ Multi-step (3+ decisions or handoffs)
- ▶ Cross-system (2+ source systems)
- ▶ Mixes structured execution with judgement



LENS 03

Data & Knowledge Fabric

Does the agent have what it needs to think?

- ▶ Structured + unstructured inputs accessible
- ▶ SOPs, policies, and history available
- ▶ Data quality fits chosen autonomy level

LENSES 4 — 6

Is it safe to ship?



LENS 04

Autonomy Spectrum

How much agency is appropriate — today and in 12 months?

- ▶ Assist · Augment · Automate · Autonomous
- ▶ Trajectory mapped, not assumed
- ▶ Human-in-the-loop placement intentional



LENS 05

Risk & Reversibility

If the agent errs, what breaks — and can we undo it?

- ▶ Blast radius mapped (financial, legal, reputational)
- ▶ Regulatory & audit exposure understood
- ▶ Reversibility & HITL controls designed in



LENS 06

Value Equation

Does the math work — and does it compound?

- ▶ Hard \$: cost-out, revenue lift, working capital
- ▶ Soft \$: cycle time, quality, talent leverage
- ▶ Scales across entities; first value < 90 days

THE AUTONOMY SPECTRUM

Pick the right level of agency.

LEVEL 1

Assist

Agent suggests; human decides and acts.

LEVEL 2

Augment

Agent drafts; human reviews and approves.

LEVEL 3

Automate

Agent acts within guardrails; humans review exceptions.

LEVEL 4

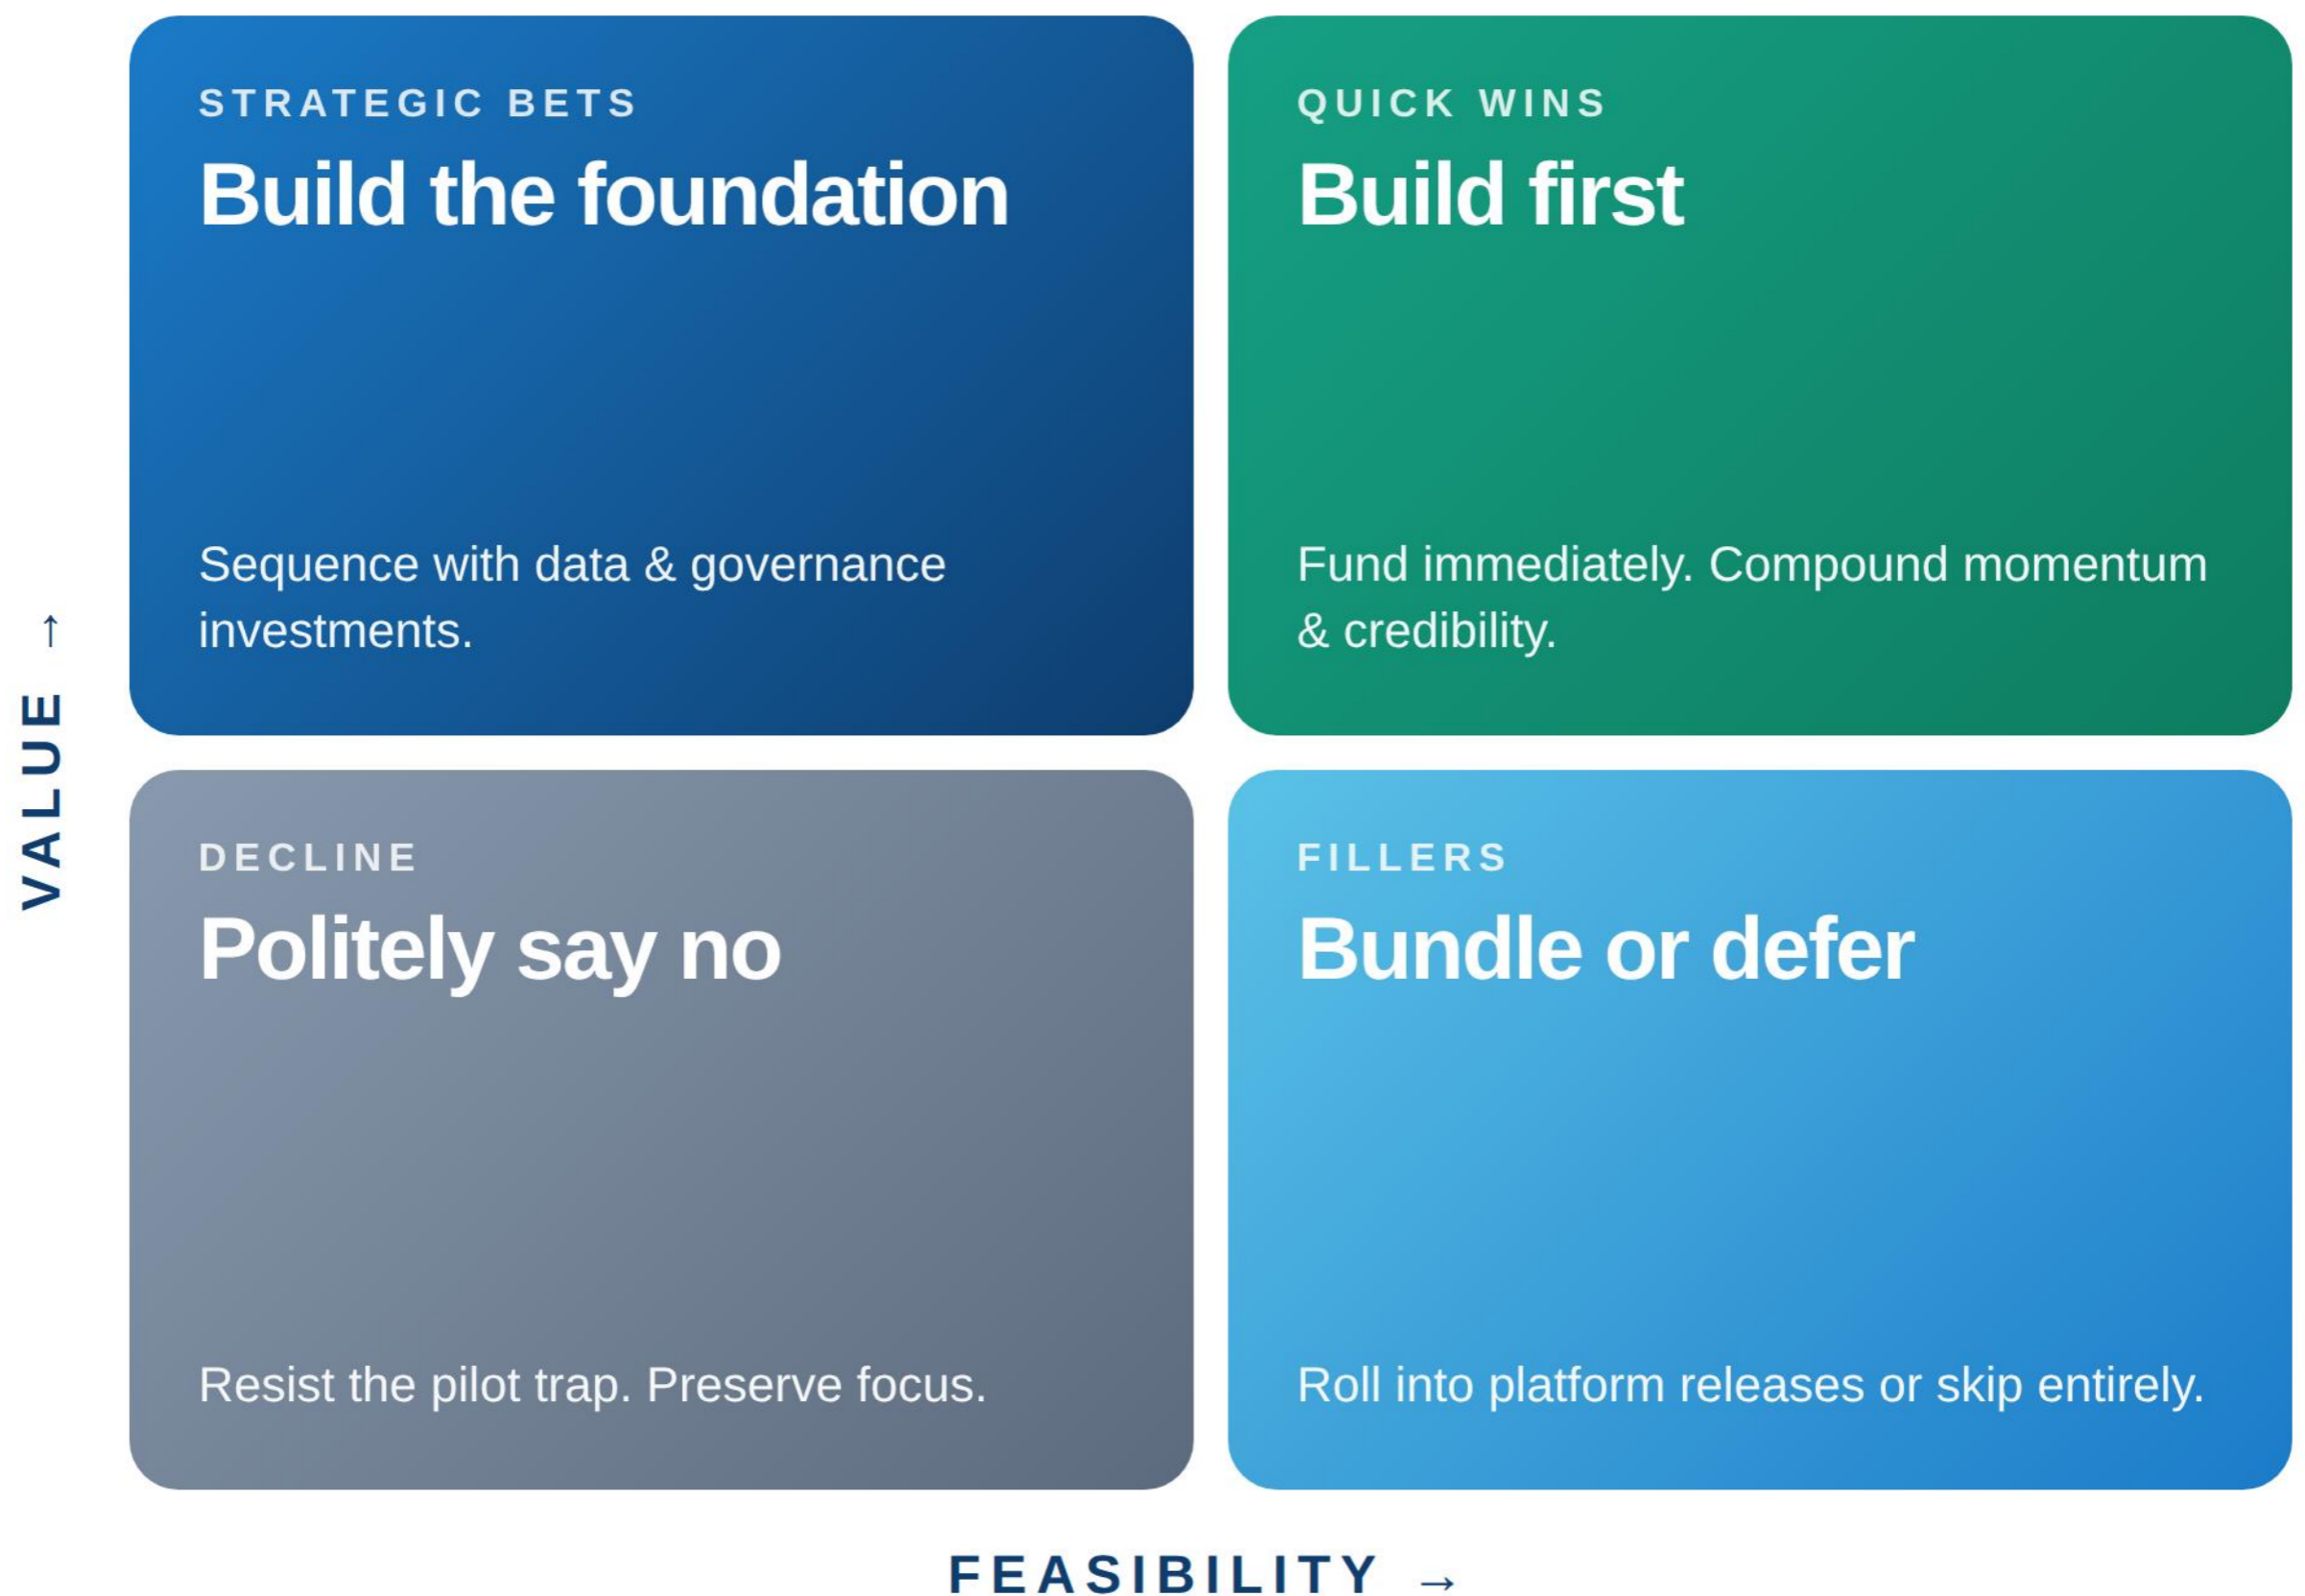
Autonomous

Agent acts; only true exceptions escalate.

Rule of thumb: Start one level lower than you think. Earn autonomy through evidence, not optimism.

PRIORITISATION

Plot. Sequence. Ship.



FINANCE

WORKED EXAMPLE

Autonomous Cash Application & Collections

O	Outcome DSO -5 to -7 days; working capital released	●●●●●
P	Process Bank match → clear → dunning → escalate	●●●●●
D	Data AR ledger, bank feeds, history — all in OneLake	●●●●●
A	Autonomy Augment now; Automate within policy	●●●●●
R	Risk Highly reversible; HITL on disputes & write-offs	●●●●●
V	Value BRC Industrial Saudia: 70%+ touchless AP proof	●●●●●

QUADRANT → **QUICK WIN**
Data is in place, risk is reversible, value is measurable. Ship in 90 days.

HR

WORKED EXAMPLE

Talent Acquisition Triage

O	Outcome Time-to-fill -40%; quality-of-hire preserved	● ● ● ● ●
P	Process JD → source → screen → shortlist → schedule	● ● ● ● ●
D	Data ATS + CVs accessible; quality variable	● ● ● ● ●
A	Autonomy Augment only — recruiters approve every shortlist	● ● ● ● ●
R	Risk ⚠ Discrimination & UAE/KSA labour-law exposure	● ● ● ● ●
V	Value High leverage; compounds across entities	● ● ● ● ●

QUADRANT → STRATEGIC BET
High value, but bias governance and HITL must precede scale. **Build the foundation first.**

SALES

WORKED EXAMPLE

Account Research & Personalised Outreach

O	Outcome Meeting acceptance x2; SDR capacity x3	●●●●●
P	Process ICP → trigger → personalise → cadence	●●●●●
D	Data CRM + public web + news — clean enough	●●●●●
A	Autonomy Augment now; Automate within brand guardrails	●●●●●
R	Risk Brand & spam exposure; fully reversible	●●●●●
V	Value Direct CAC impact; first value < 90 days	●●●●●

QUADRANT → **QUICK WIN**
Self-contained, fast time-to-value, brand-guardrailed. A textbook first build.

BEFORE YOU FUND ANYTHING

The Three- Question Gate.

1

Can a human SME describe success in one sentence?

If not, the outcome isn't crisp enough to engineer.

2

Can we instrument the outcome?

If we can't measure it weekly, we can't improve it.

3

Are we willing to cede autonomy at the chosen level?

If governance won't allow it, scope down or stop.

Any "No" → not yet ready. Fix the gap before funding the build.

THE PATTERN

Quick Wins fund Strategic Bets.

Don't start with the moonshot.

Start with the use case that ships in **90 days**, generates measurable \$, and earns you the right to attempt the harder ones.

F-OS™ · AGENTIC PROCESS OUTSOURCING

**We sell work, not software.
We deploy agents, not tools.**

Multi-agent platform for the Office of Finance, orchestrated by HOFT's human SMEs as the human-in-the-loop layer.

What's your team scoring next quarter?

DM me — let's pressure-test it together.