Build agentic AI around real commercial work

NSG helps UK SMB and mid-market businesses identify where agentic AI can improve workflow performance, reduce manual drag and strengthen commercial control.
We assess the process, data, governance and commercial case before recommending where AI agents should assist, automate or stay well away.

AGENTIC AI IS NOT A BETTER CHATBOT. IT IS A WORKFLOW DECISION.

Agentic AI is becoming one of the most talked-about areas in business technology.

Predictably, that means it is also being misunderstood at speed.

Many leaders are now being told that AI agents will take over admin, automate operations, run workflows, support customers, manage internal tasks and free teams from repetitive work. Some of that is true. Some of it is sales theatre wearing a technical jacket. The difference depends on how the agent is designed, what it is allowed to do, which systems it can access, how its work is checked, and whether the underlying business process is fit for automation in the first place.

An AI agent is not just a chatbot with a better job title.

A chatbot responds.

An agent acts.

That distinction is important. Traditional AI tools usually wait for a user to ask a question or give an instruction. They generate a response, produce a draft, summarise information or answer based on available context. Agentic AI is different because it is designed to work towards a goal. It can break work into steps, use tools, interact with systems, check progress, make adjustments and continue moving through a task.

I

Where agentic AI can create commercial value

Agentic AI becomes useful when work is repetitive, multi-step, rules-based enough to control, and currently too manual for the value it creates. That does not mean every process should be automated.

Some work needs judgement. Some work needs empathy. Some work needs commercial experience. Some work needs human negotiation, political awareness or contextual interpretation. Trying to automate those areas too early is how businesses create expensive mistakes while congratulating themselves for being innovative.

The better starting point is work that already has a recognisable process.

Examples include:
Customer support triage, internal helpdesk routing, meeting follow-up workflows, document intake, proposal preparation, CRM data enrichment, onboarding tasks, policy Q&A, invoice query handling, service request classification, compliance evidence gathering, renewal administration, sales follow-up preparation and management reporting.

These are not glamorous use cases. That is why they are useful.

The best early AI use cases are often boring. Boring means repeatable. Repeatable means controllable. Controllable means measurable. Measurable means the business can decide whether the agent is actually creating value or just producing more digital noise.

Why broken processes should not be automated

Agentic AI does not rescue bad process design. It exposes it.

If the business does not have clear workflows, clean data, defined approval points, documented rules or accountable owners, an AI agent will struggle. Worse, it may appear to work while quietly multiplying errors across the business.

This is the danger.

A manual process with poor controls is inefficient. An automated process with poor controls is dangerous.

Before deploying agentic AI, leaders should ask some uncomfortable questions:

Is the workflow already understood?

Are the decision rules clear?

Is the data reliable?

Are exceptions known?

Who owns the outcome?

What should the agent never be allowed to do?

Where does human approval become mandatory?

How will we know if the agent made a mistake?

How will we audit what happened?

If these questions cannot be answered, the business is not ready for autonomous action. It may still be ready for assisted AI, but not for agents that execute work across systems.

There is a difference between AI helping people work better and AI being allowed to move work forward.

That difference needs governance.

Agentic AI and Commercial Transformation

For NSG, agentic AI sits inside the wider Commercial Transformation agenda.

Commercial Transformation is about improving how a business converts effort into growth, control and operational performance. That includes GTM execution, revenue operations, customer handoffs, data flow, management rhythm, automation and decision support.

Agentic AI can support all of that, but only if it is tied to real business outcomes.

The wrong approach is to start with:
“We need AI agents.”

The right approach is to start with:
“Where does work get stuck, repeated, delayed, misrouted, forgotten or poorly handed over?”

That is the commercial lens.

An AI agent should be designed around a defined friction point.

For example:

A lead enters the business but is not qualified quickly enough.

A sales follow-up is delayed because discovery notes are messy.

A proposal takes too long because information sits across calls, emails and documents.

A customer issue is passed between teams because ownership is unclear.

A manager cannot trust forecast commentary because CRM updates are inconsistent.

An operations team spends hours chasing basic information.

A finance team manually reconciles queries that follow known patterns.

These are agentic AI opportunities because they involve process, information, rules and action.

They are not just AI opportunities.

They are operating model opportunities.

Agentic AI and Revenue Operational Control

Revenue Operational Control is one of the clearest areas where agentic AI can help.

Most growing businesses do not lose control of revenue because people are lazy. They lose control because the operating system around revenue is too informal.

Pipeline definitions are unclear. Qualification standards vary. CRM hygiene is poor. Handoffs between marketing, sales, delivery and finance are inconsistent. Forecasting depends too much on opinion. Follow-up actions are missed. Deal risks are not surfaced early enough. Customer information is scattered across inboxes, calls, documents and personal notes.

Agentic AI can help tighten this.

It can support revenue teams by:
structuring discovery notes, identifying missing qualification data, summarising deal risk, preparing follow-up actions, checking CRM completeness, prompting next steps, drafting customer communications, reviewing handoff quality, flagging stalled opportunities and supporting pipeline governance.

But it must sit inside a proper revenue operating model.

If the business has not defined what a qualified opportunity is, the agent cannot enforce it properly.

If the business has no agreed sales stages, the agent cannot reliably support stage progression.
If managers do not inspect pipeline quality, the agent becomes another ignored system prompt.
If CRM data is weak, the agent will simply reason from weak inputs.

Agentic AI can strengthen revenue control.
It cannot replace revenue discipline.

That is the bit people often skip, usually because discipline is less exciting than pretending software has solved management.

Agentic AI and Growth GTM

There is also a strong connection between agentic AI and Growth GTM.

Modern GTM execution involves many moving parts: market segmentation, account research, contact enrichment, message development, outreach sequencing, engagement tracking, lead qualification, campaign reporting and sales handoff.
Much of that work is structured, repeatable and data-heavy.

Agentic AI can support GTM by helping teams: research target accounts, identify trigger events, summarise market signals, prepare account briefs, generate message variants, classify responses, recommend follow-up actions, update campaign records and provide performance summaries.

However, this does not mean businesses should automate their entire GTM engine and hope buyers respond positively to yet more robotic outreach. Humanity has endured enough lazy email sequencing already.

The value is not in sending more.

The value is in improving relevance, timing, consistency and follow-through.

Agentic AI should help GTM teams become more precise, not more noisy.

That means the agent must operate within clear commercial rules:
who the business targets, why that market matters, what pain is being addressed, what evidence supports the approach, what messages are approved, what claims are allowed, when sales should intervene and how engagement should be measured.

Without that structure, agentic AI risks becoming a faster way to annoy more people.

Human oversight is not optional

A useful way to think about an AI agent is as a junior operator.

It may be fast. It may be capable. It may handle repetitive tasks well. But it still needs instructions, limits, supervision and review.

No sensible business would hire a junior employee on Monday, give them access to finance systems, customer data, sales records and legal documents on Tuesday, then let them act without oversight by Wednesday.

Yet some businesses seem oddly willing to do the AI version of that.

Agentic AI needs clear authority levels.

Some actions can be fully automated.

Some actions should be drafted for human review.

Some actions should only be recommended.

Some actions should be blocked entirely.

For example, an agent may be allowed to summarise a customer issue automatically. It may be allowed to draft a response. It may be allowed to create an internal task. But it may need human approval before sending anything externally, changing commercial terms, updating contractual records, approving refunds, modifying finance data or making decisions with legal or compliance implications.

This is basic control.

Not resistance to innovation.

Control is what allows innovation to scale without turning the business into a live experiment.

The practical adoption path

The best adoption path for agentic AI is staged.

Start with low-risk assistance. Move into workflow support. Then consider controlled automation. Only later should the business explore higher-autonomy agents.

A sensible path looks like this:

First, use AI for knowledge retrieval, summarisation and drafting where people remain fully in control.

Second, introduce guided workflow support, where AI prepares actions, checks completeness and recommends next steps.

Third, allow agents to trigger low-risk internal actions, such as creating tasks, routing tickets, updating non-sensitive records or producing standard reports.

Fourth, test higher-autonomy workflows in controlled areas with strong logging, escalation and review.

Fifth, only scale once performance, governance, cost and risk are properly understood.

This progression matters because autonomy should be earned.

An agent should not be trusted because it sounds impressive. It should be trusted because it has performed reliably within a known process, under defined controls, with measurable results.

What leaders should assess before deploying agents

Before investing in agentic AI, leaders should assess six areas.

The first is workflow readiness. Is the process clear enough to automate or assist?

The second is data quality. Can the agent access reliable, current and approved information?

The third is system access. Which tools does the agent need, and what permissions should it have?

The fourth is risk level. What happens if the agent gets it wrong?

The fifth is human oversight. Where does approval, review or escalation sit?

The sixth is commercial value. What measurable improvement is expected?

This last point matters.

Agentic AI has cost. It may require workflow design, integration, licences, API usage, monitoring, testing, governance and ongoing improvement. Multi-step agents can also consume more compute and operational resource than simple AI interactions.

So the business case needs to be specific.

Will it reduce response time?

Will it reduce admin effort?

Will it improve data quality?

Will it increase follow-up consistency?

Will it reduce handoff failure?

Will it improve compliance evidence?

Will it reduce cost to serve?

Will it improve conversion through better timing and control?

If the answer is vague, the use case is not ready.

Shadow AI and Growth GTM

Shadow AI also affects Growth GTM.

Marketing and GTM teams are under constant pressure to produce content, campaigns, messaging, insight, research, segmentation and outreach at speed. AI is obviously useful here.

But unmanaged AI use can create serious problems. Generic content. Weak differentiation. Unsupported claims. Poor targeting. Incorrect market assumptions. Over-personalised outreach. Misuse of prospect data. Brand inconsistency. Duplicate content. Poor source control. Campaign volume without commercial relevance.

AI can make GTM faster.

It can also make bad GTM faster.

That is the issue.

The aim should not be to use AI to produce more noise. The aim should be to use AI to improve market focus, buyer relevance, message quality and campaign discipline.

That requires controlled source material, approved messaging architecture, clear claim boundaries, proper review, and a direct link between campaign output and sales follow-up.

Shadow AI breaks that link.

Governed AI strengthens it.

How leaders should bring Shadow AI under control

The practical answer is not complicated, but it does require discipline.

First, discover what is already happening.
Do not assume staff are waiting for permission. They are probably not. Run an AI usage review. Ask teams what they are using, what they are using it for, what problems it helps solve and what data they enter. Make it safe for people to be honest. If the review feels like a disciplinary trap, people will hide the truth, because apparently self-preservation remains popular.

Second, classify the use cases.
Separate low-risk drafting and productivity use from higher-risk use involving personal data, customer information, contracts, financial information, HR records, regulated content, automated decisions or external communication.
Third, define approved tools. Give employees clear options. If the business wants people to stop using random public tools, it must provide approved alternatives that are easy to access and good enough to use.

Fourth, create plain-English rules.
Policies should explain what people can do, what they cannot do, what data must never be entered, what needs review and where to ask for help. Role-based examples matter more than abstract principles.

Fifth, fix permissions and data access.
Before scaling AI across Microsoft 365 or other internal systems, review access rights, oversharing, sensitive data exposure, retention, labelling and document governance. AI will surface what your permission model already allows.

Sixth, monitor without making work impossible.
Security monitoring should identify risky tools, high-risk access patterns and unusual activity. But it should not create so much friction that people go back to hidden behaviour.

NSG view: agents should be designed around control, not novelty

At Noodle Spark Group, our position is straightforward.

Agentic AI should not be introduced because it is new.

It should be introduced where it improves the way work flows through the business.

That means starting with business friction, not technical possibility.

The strongest use cases usually sit where teams are wasting time, repeating structured tasks, relying on scattered information, missing handoffs, producing inconsistent outputs or failing to act quickly enough on known signals.

That is where agents can help.

But the agent must sit inside a designed operating model. That includes governance, process ownership, system integration, data controls, human review, audit trails and clear success measures.

Without that, agentic AI becomes another unmanaged technology layer.

With it, agentic AI can become a practical part of Commercial Transformation.

It can reduce drag.

It can improve consistency.

It can strengthen revenue control.

It can support better GTM execution.

It can help smaller teams operate with more discipline and capacity.

But only if leaders stop treating AI agents as digital magic and start treating them as controlled operational workers with defined authority.

That is the grown-up version of agentic AI.

Less exciting on LinkedIn.

Much more useful in the real world.

STRONGER TOGETHER