A pragmatic guide for business leaders who want to cut through AI hype, understand what today's tools are good at, and make confident, value-driven decisions without needing to become AI experts.

AI Without the Hype

What it's good at (and not)

Morgan Stern, Keryk_AI

Introduction: The AI Hangover Is Here

Over the past two years, AI has gone from something discussed in strategy decks to something that executives are expected to be implementing. Whether its GenAI, LLMs, autonomous agents, chatbots or copilots, everywhere you turn, there's a new product or platform promising to change the way you work.

And yet, for all the noise, most companies are still stuck in the same place: trying to figure out where to start, what's real, and what might drive actual business value.

The truth is you don't need to be an "AI expert" to make smart decisions about AI. You just need to be pragmatic. Don't get distracted by the mechanics. Don't feel pressured to follow hype-driven trends. And above all, don't assume that the presence of an AI model automatically makes something valuable.

This paper is written for executives, curious mid-level managers, and anyone tasked with figuring out how AI fits into the bigger picture. We'll cut through the noise and help you focus on the only question that matters:

Can this technology consistently help you get the output you need, at a cost that makes sense?

If it can, great. If not, it's probably not worth the effort right now, no matter how flashy the demo.

The Hype Cycle in Overdrive

AI is no stranger to hype. But over the last 18 months, the cycle has gone into overdrive. You've likely heard things like:

- "LLMs can reason like a human."
- "Agents will automate entire workflows."
- "AI will replace 80% of jobs."

Behind many of these claims is a well-polished demo, a funding round, or a vendor with a new framework. But here's the hard truth: **most of these promises collapse under the weight of real-world complexity.**

The problem isn't just the technology, it's the context.

Real-World AI Misfires

Let's look at a few cautionary tales:

- Zillow's AI-Powered Home Pricing Model

  Billed as a sophisticated way to buy and flip homes at scale, it relied heavily on predictive models that didn't account for market volatility. The result? Hundreds of millions in losses and a total shutdown of the initiative.

- IBM Watson for Oncology

  Promoted as a revolutionary tool for cancer treatment recommendations, it struggled with accuracy, adaptability, and practical integration with clinician workflows. Despite years of marketing, the project never delivered the value it promised.

These weren't fringe use cases. These were high-budget, high-stakes initiatives. And they failed not because AI doesn't "work", but because expectations were disconnected from reality.

The Agent Hype

One of the most inflated bubbles today is around **agentic frameworks**, the idea that you can spin up AI "agents" to autonomously plan, reason, take actions, and learn. It sounds futuristic. But in practice?

- Many agents today are just choreographed sequences of API calls and LLM prompts.
- They require constant oversight and break easily in unstructured environments.
- Their "reasoning" often amounts to surface-level prediction, not deep understanding.

That doesn't mean agents are useless. It just means they're not magic. (I personally am bullish on agents and agentic-frameworks and have seen great advances in the past couple months with releases like the OpenAI Agents SDK and discussions around A2A)   However, If you build your business plans around what they *might* be able to do one day, you're likely to waste time, money, and trust along the way.

The Reality Gap

The current gap is twofold:

- The delta between what AI tools are capable of and what your business actually needs
- The delta between what a model can do in theory and what your organization can reliably deploy, monitor, and support in production

And in between those two gaps? A lot of expensive experiments that don't deliver.

The AI Toolbox: What's Actually Out There

When people talk about "AI," they're often referring to large language models (LLMs) like ChatGPT. But AI is a much broader field, and many of the most useful, battle-tested tools aren't the ones making headlines.

Let's take a step back and look at a more well-rounded toolbox.

1. Generative AI (GenAI) & Large Language Models (LLMs)

- **What it does:** Generates human-like text, code, summaries, and responses.
- **Strengths:** Excellent for drafting, rewording, summarizing, ideating.
- **Limitations:** Outputs are variable. It does not *know* anything, it predicts likely next words.
- **Best use cases:** Internal drafting tools, customer service triage, rapid experimentation.

2. Traditional Machine Learning (ML)

- **What it does:** Learns patterns from historical data to make predictions or classifications.
- **Strengths:** Highly accurate when trained on quality data.
- **Limitations:** Requires clean, labeled data and consistent patterns.
- **Best use cases:** Forecasting demand, risk scoring, customer churn prediction, fraud detection.

3. Computer Vision

- **What it does:** Analyzes images or video to detect objects, classify scenes, or measure change.
- **Strengths:** Great for visual inspection, compliance checks, and physical-world automation.
- **Limitations:** Needs large image datasets, careful tuning, and clear visuals.
- **Best use cases:** Manufacturing QC, safety monitoring, logistics and infrastructure inspection.

4. Optimization & Operations Research Models

- **What it does:** Finds the most efficient way to allocate resources, schedule tasks, or route flows.
- **Strengths:** Excellent for structured, rule-based decision making.
- **Limitations:** Requires clearly defined goals, constraints, and parameters.
- **Best use cases:** Workforce planning, logistics routing, equipment utilization, network optimization.

5. Agents and Autonomous Systems (The Shiny New Thing)

- **What it does:** Sequences LLMs and tools to attempt multi-step tasks.
- **Strengths:** Exciting potential for delegation of complex, multi-step processes.

- **Limitations:** Very fragile, high hallucination risk, still immature for production use.
- **Best use cases (today):** Internal experimentation, sandboxed environments, co-pilot-style tasks with human supervision.

Use the Right Tool for the Job

Here's the bottom line:

LLMs are not the only game in town. And in many cases, they're not the best tool for the job.

While GenAI can be powerful, more grounded and mature technologies like predictive models, rule-based systems, or optimization algorithms often deliver better, more consistent value in critical environments.

Choosing the right tool means asking the right questions.

The Problem with Non-Determinism

One of the most misunderstood characteristics of LLMs is that they are **non-deterministic** by design.

That means: The same input doesn't always produce the same output.

This might be acceptable for creative tasks like writing, brainstorming, or ideating. But for many business processes, **consistency is non-negotiable**.

LLMs Are Improv Artists, Not Calculators

LLMs work by predicting the most likely next word or token based on the prompt and their training data. They don't "understand" context the way a person does. Instead, they're generating plausible responses based on patterns.

That's why:

- They sometimes make things up (hallucinations).
- They can give different answers to the same question.
- They occasionally offer confident answers that are completely wrong.

If your use case requires **precision, consistency, traceability, or compliance**, that's a serious red flag.

When Variability Becomes a Liability

Let's look at a few example scenarios:

Use Case	Is Variability Acceptable?
Summarizing internal meeting notes	✅ Yes
Drafting marketing headlines	✅ Yes
Answering HR policy questions to employees	❌ No: must be accurate, consistent
Generating billing or pricing data	❌ No: must be exact

Guardrails Are Not a Cure-All

Vendors often try to patch over these issues with tools like:

- Prompt engineering frameworks
- Guardrails and filters
- Role-based constraints

These can help. But if the core tech wasn't built for consistency, you're fighting the current. In some cases, it's a sign you're using the wrong tool for the problem.

Ask Yourself:

- What happens if this output is occasionally wrong?
- Can I live with variation or does it undermine the process?
- Am I spending more time building controls than I would with a simpler tool?

AI doesn't need to be perfect but in many business processes, it does need to be consistent, auditable, and right most of the time.

When it's not? It's time to rethink the approach.

Framework: Is This a Good Use Case for AI?

Let's simplify the decision-making.

Not every task needs AI. Not every business problem benefits from it. And when AI is used, not every type of AI delivers the same value. So how do you know if your use case is a good fit?

Start by asking five simple but revealing questions.

1. Does the AI Technology Natively Support the Business Need?

- Is this problem something AI is naturally good at solving?
  (Pattern recognition? Text generation? Forecasting?)
- Or are you bending the problem to fit the tech?

If you find yourself building layer after layer of prompts, filters, retries, and wrappers just to keep things stable... that's a red flag.

2. How Much Variability is Acceptable?

- Can the process tolerate different outputs each time?
- If the answer isn't always the same, is that a feature or a failure?

Creative tasks = high tolerance.

Compliance-driven tasks = low to zero tolerance.

3. What Are the Consequences If It's Wrong?

- Will someone catch a bad result before it matters?
- Is there risk to the customer, employee, or bottom line?

Always match the risk of error to the maturity of the solution.

4. Does the Cost Make Sense for the Business Value?

- Some models are expensive to run, especially LLMs with large context windows.
- Are you using $50 worth of tokens to answer a $5 question?

Value comes from impact, not novelty.

🔍

5. Can the Output Be Easily Measured and Trusted?

- Can you track quality, performance, and improvement over time?
- Or is the system a black box that nobody wants to challenge?

If people can't trust the results, they won't use them, no matter how "intelligent" it is.

The Guiding Principle: Output First, Tech Second

AI should be treated like any other tool. The only question that matters at the end of the day is:

Does it consistently produce an output that meets the business need in a way that is better, faster, or cheaper than what came before?

If yes, great. If not, move on or explore alternative solutions.

What to Do Instead

So now that we've walked through the hype, the limits, and the right questions to ask… what *should* a business leader do to make smart decisions about AI?

It comes down to three simple things:

1. Focus on the **business need**
2. Evaluate outputs, not mechanics
3. Apply critical thinking, not blind faith

Let's break that down.

1. Start with the Outcome You Want

Before looking at any technology, ask:

- What result are we trying to achieve?
- What's slowing us down today?
- Where do we spend time, money, or effort that we'd love to reduce?

Then ask:

"Could any tool (AI or non-AI)l help us get there more efficiently, more accurately, or more reliably?"

If the answer is yes, great. If not, don't force it.

2. Evaluate Tools by the Output They Deliver

It doesn't matter how cool the underlying model is. What matters is:

- Does it work reliably in your context?
- Can it deliver value without major oversight?
- Are the results consistent enough for how you want to use them?

You don't need to understand the math behind a hammer to know it drives nails. Same with AI—if it gets the job done better, use it. If it doesn't, shelf it.

3. Critical Thinking is Your Competitive Advantage

With so much AI noise in the market, the best thing you can do is apply sound judgment:

- Be skeptical of tools that require extreme customization to fit your problem.
- Watch for escalating costs in token usage, retraining, or infrastructure.
- Keep asking, "Does this actually solve a real problem for us?"

Real Talk: Some Use Cases Just Aren't Worth It (Yet)

Yes, you *can* build an agent that uses an LLM to schedule meetings, extract action items, and summarize CRM entries.

But if it needs 14 layers of prompt engineering, three APIs, constant human review, and on-going care and feeding, maybe a smart intern is a better fit for now.

AI should simplify your business, not overcomplicate it.

Conclusion: You Don't Need to Be an AI Expert; You Just Need to Be Outcome-Oriented

The goal of this guide wasn't to make you fluent in AI jargon. It was to give you the perspective and confidence you need to make smart, value-driven decisions.

You don't need to know how an LLM works under the hood.

You don't need to hire a fleet of AI engineers.

You don't even need a chatbot (unless it actually solves something for your team).

What you *do* need is clarity:

What outcome are you after?

Can AI help you get there, consistently and cost-effectively?

And if not, what other path might be smarter?

Stay focused on **outputs, value, and fit** and you'll already be ahead of the pack.

Ready to Get Started the Right Way?

If you're serious about cutting through the hype and finding *practical, business-ready opportunities* to use AI, we'd love to help.

- **Keryk's Jumpstart Program** is designed for exactly that:
  A short, high-impact engagement that educates your team on the AI opportunities and identifies strong use cases for your business, with clear next steps.
- **Want to talk first?**
  Let's have a quick call. We'll help you sort out what's signal and what's noise - no pressure, no jargon, just clarity.

Visit https://www.keryk.ai to book your **Jumpstart** or schedule a **free intro call**.

Visit **keryk.ai** for more insights and offerings.