

**everyone
is smart
now.**

Everyone Is Smart Now.

AI is making intelligence abundant. The interesting question is not how much more we can make with it, but what becomes valuable when making is no longer the hard part.

In 1968, an Andy Warhol exhibition catalogue carried the prediction that everybody would be world-famous for fifteen minutes. The internet eventually performed a peculiar variation on the idea. It gave almost everybody access to the machinery of publicity while leaving attention stubbornly scarce.

Publishing became easy; being read did not. Recorded music became almost universally accessible; our capacity to listen did not expand with the catalogue. Billions of capable cameras did not produce billions of memorable photographs. Each technology removed a constraint, and each time the economic or cultural value migrated towards something the technology could not make abundant.

The difficult part moved.

Something similar is happening with intelligence.

A founder can describe an idea in ordinary language and have working software remarkably quickly. A marketing team can explore campaigns while the brief is still changing. An analyst can interrogate a market in minutes. Inside companies, people who would once have waited in an engineering queue can increasingly make useful software themselves.

The obvious story is that machines are becoming smarter. The more consequential one is that useful cognitive capability is becoming cheaper and more widely available.

I have spent much of my career watching technologies cross versions of this threshold. I began as a software engineer when making software required relatively scarce technical skills. At Ford, technology was changing the systems through which physical products were engineered and manufactured. At Microsoft, I lived through the transition from software towards cloud and AI, as computing infrastructure companies had once owned increasingly became capacity they could consume. More recently, at Amazon, generative AI began changing the economics of creative production.

These transitions were different in almost every practical respect, but they shared a pattern that is easy to miss while a technology is new. When a previously scarce capability becomes widely available, the capability still matters; it simply stops explaining as much of the difference between winners and everyone else.

Value moves towards the constraint that remains.

AI is beginning to move several of them at once.

Plausibility Gets Cheap

For decades, a prototype carried information because building one was difficult. Somebody had assembled enough expertise, capital and conviction to drag an idea into

reality, so the existence of the object told you something about the seriousness of the undertaking.

That distance between idea and artefact used to be substantial. At Ford, changing how vehicles were designed and manufactured meant confronting engineering systems, physical production and organisational dependencies that did not yield because an idea was persuasive. Later, in enterprise technology, even digital propositions could require months of architecture, engineering and negotiation before anybody saw something convincing.

Difficulty was inefficient, but it also performed an accidental filtering function. Not every idea survived long enough to become tangible.

AI weakens that filter.

An idea can now acquire the surface characteristics of a product before much external reality has touched it. Describe an application, iterate through conversation, add a credible interface and something that barely existed yesterday can arrive in a meeting looking surprisingly complete.

None of this means that durable products have become trivial to build. Security, architecture, integration, reliability, regulation, maintenance and economics remain stubbornly real. What has changed is the cost of producing something plausible enough to create belief before those questions have been answered.

AI has industrialised plausibility.

That matters because tangible ideas behave differently inside organisations. Before the prototype exists, people debate whether the idea deserves to exist. Once they can see and use it, the conversation tends to move forward. Features are proposed, adjacencies appear, a market size is calculated and the thing begins accumulating a future.

The prototype starts lobbying for its own survival.

This is not an argument against prototyping; cheaper experimentation is one of AI's genuine advantages. The distinction is between using a prototype to learn and allowing its existence to masquerade as evidence that the underlying business deserves commitment.

When prototypes were expensive, producing one demonstrated a degree of capability and conviction. As convincing prototypes become cheaper, technical feasibility carries less information about commercial consequence. The danger is therefore not merely a proliferation of bad technology, because bad technology is often straightforward to reject. It is the proliferation of competent technology attached to weak propositions: applications nobody changes behaviour for, agents improving processes that should perhaps disappear, and products whose customers remain theoretical.

The landfill of the AI age may be filled with things that work.

Reality Becomes the Proof

If plausibility becomes cheaper, evidence has to move somewhere harder to manufacture.

A customer changing behaviour is harder to generate than a prototype. So is a buyer reallocating budget, a sales process becoming repeatable, retention surviving novelty, a channel producing demand economically or margins improving as volume grows. These are imperfect signals—business rarely offers laboratory conditions—but they contain something a demonstration cannot: consequence outside the imagination of the people building it.

Someone has surrendered something scarce.

That might be money, time, attention, data, reputation or an established habit. Whatever its form, the surrender tells you that the proposition has encountered reality and altered it.

Working across Microsoft's businesses internationally made this distinction difficult to ignore. The same underlying technologies encountered healthcare, banking, manufacturing, retail and government very differently because those industries operated differently, bought differently and carried different constraints. Even across markets, a proposition that travelled effortlessly as technology could encounter completely different buying behaviour, regulation, distribution and economics.

Technology travels more easily than behaviour.

That suggests a reversal in the conventional logic of product development. Instead of using the prototype to prove the market, use the market to prototype the proposition. Test price while the product is still fluid. Discover whether somebody will allocate budget before building the organisation required to serve them. Understand what customers must stop doing, not merely what they say they would like to start doing. Test the route to demand while changing course remains cheap.

This is where the distinction between product and go-to-market begins to dissolve. If technology can be altered continuously, proposition, price, distribution, customer behaviour and economics can be discovered as parts of the same system. The point is not to commercialise something after it has been built, but to allow commercial reality to shape what gets built in the first place.

A product without a credible route to demand is not finished. It is merely built.

Abundance Meets a Human Being

The internet gave us something close to an infinite shelf. AI is placing an extraordinarily productive factory behind it.

Companies can produce more software, imagery, propositions, campaigns and services than their customers could reasonably consume. The supply of competent things can expand dramatically without any corresponding increase in human attention, organisational capacity or purchasing power.

Nobody gets a twenty-fifth hour.

This is why cheaper production can make distribution more, rather than less, consequential. A company capable of generating endless propositions but unable to reach and retain customers economically has not solved its growth problem. It has expanded the number of ways in which it can discover that it has one.

I saw a version of this at Amazon's Brand Innovation Lab. Generative AI could increase the speed and range of creative possibilities available to global brands, but those brands rarely suffered from a shortage of things they *could* say. The harder question was which idea belonged in which context, for which customer, at which moment, and whether it changed anything that mattered.

More creative supply does not manufacture more human attention.

The same logic extends beyond advertising. As functionality becomes easier to reproduce, established relationships, distribution, ecosystems and positions inside important workflows can become more consequential. Agents may move this boundary again if they increasingly act across services on a user's behalf. Digital companies spent decades trying to own destinations because customers had to visit their websites, apps and portals to accomplish something; if some intentions can instead be executed by intermediating systems, parts of that interface may shift.

How far it shifts remains uncertain. The useful strategic question is not to pretend we know, but to ask what would remain scarce if it did.

Usually the answers are less glamorous than intelligence: a trusted relationship, permission to act, knowledge of the current situation, access to demand.

The Company Knows Tuesday Morning

This is also why saying “proprietary data is the moat” is too crude.

A capable general model can know an extraordinary amount about banking without knowing why a particular payment failed five minutes ago. It can understand manufacturing without knowing which component failed inspection this morning. It can absorb libraries of sales methodology without knowing that an important customer began evaluating a competitor yesterday.

General intelligence can know a great deal about the world. Businesses operate in situations.

What matters in those situations is often context: customer history, workflow state, relationships, permissions, operational reality and current intent assembled at the moment a consequential decision has to be made. Some of that context can compound as interactions produce knowledge, knowledge improves decisions and better outcomes deepen relationships.

The underlying intelligence may increasingly be available from several providers. The accumulated position from which a company applies that intelligence can be much harder to reproduce.

This does not mean context automatically becomes a moat; very little automatically becomes a moat. Customers can leave, data can be portable, competitors can acquire similar access and platforms can change the economics underneath you. But it suggests that defensibility in AI may depend less on possessing intelligence in the abstract and more on embedding it somewhere consequential.

Where is the system? What can it see? What happened five minutes ago? What is it permitted to do next?

Those are business questions disguised as technical ones.

Creation Is Cheap. Inheritance Isn't.

Inside the organisation, abundance creates a different problem.

Most large companies have some version of the spreadsheet nobody wants to delete. It began as a useful model, accumulated tabs, formulas and dependencies, lost its original author and eventually became important enough that nobody wants to discover what happens if a particular cell changes. Spreadsheet risk is well documented; the joke persists because the organisational pattern is familiar.

AI can extend the pattern far beyond spreadsheets. Give people the ability to create software and the immediate effects can be valuable: sales builds a tool, finance automates reporting, marketing creates an agent, operations removes an irritating manual step.

Then some of the experiments survive.

People depend on them. Data enters them. Permissions accumulate. Systems connect. Their creators change jobs. Somewhere else, another team solves a similar problem differently because making something has become easier than discovering, adapting or governing what already exists.

Software can become cheap to create long before it becomes cheap to inherit.

Once something becomes consequential, it acquires obligations—security, maintenance, integration, compliance, continuity and accountability—and none disappears because version one took an afternoon. This creates a possibility leaders should take seriously without exaggerating it into inevitability: AI can improve productivity inside individual teams while allowing complexity to accumulate across the enterprise.

The relevant measure is therefore not simply whether people can produce more. It is whether the business creates more customer and economic value without complexity, cost and coordination expanding alongside it.

That requires a distinction between freedom to experiment and the right to become permanent. Rebuilding every old approval gate would squander much of what makes the technology useful, but allowing every successful experiment to harden into infrastructure creates another form of waste.

Let creation be easy. Make permanence earn its place.

An AI-native organisation may need to become as accomplished at removing things as it becomes at making them.

Permission Has a Price

There is another constraint that more intelligence does not automatically remove.

An incorrect paragraph is irritating. A system moving money, changing a production schedule, altering a customer's account or making a regulated decision is different because the issue is no longer merely whether the system knows enough.

It is whether anyone trusts it with authority.

Human organisations have always separated expertise from agency. Knowing how to drive does not give somebody the keys to every vehicle; understanding markets does not give every analyst authority over a firm's balance sheet. AI makes this distinction increasingly important as systems move from producing recommendations towards taking actions.

Trust, in that sense, has economic consequences. Systems that are permitted to act can occupy deeper workflows and perform more consequential work. Reliability, security, provenance, accountability and reversibility are therefore not simply constraints around AI; they help determine how much responsibility customers and institutions are willing to delegate to it.

The most capable system does not necessarily capture the most value. Capability has to cross the boundary into permission.

The Economics of Choosing

In 2002, David Bowie told *The New York Times* that music would become “like running water or electricity.” The interesting part of the observation was not that music would cease to matter. It was that ubiquity would alter where value could be created and captured.

Recorded music remained culturally important while its economics changed around it.

AI presents business with a related question. If knowledge, analysis, software and creative production become cheaper and more widely available, the ability to produce them does not become worthless. It simply becomes less sufficient.

Eventually this lands on something business has often treated as softer than strategy: taste.

Creative disciplines have long understood the difference between technique and taste. Technique expands what can be made; taste helps determine what deserves attention. In a company, the equivalent is judgement under abundance: recognising which problem matters, which customer behaviour constitutes meaningful evidence, which capability changes the economics, which uncertainty deserves testing next and which attractive idea should die before the organisation becomes committed to wanting it to succeed.

I have seen enough transformations to distrust the assumption that more activity necessarily means more progress. Building Microsoft's transformation business to substantial scale was not simply an exercise in generating more propositions. Across industries and markets, the harder work was distinguishing where technology genuinely changed a product, customer experience or operating model from where technology was simply looking for somewhere to go. Years later at Amazon, generative AI expanded the creative possibility space enormously without removing the obligation to choose.

This is where cheap experimentation becomes genuinely powerful. It can lower the cost of being wrong, provided companies are willing to discover that they are wrong while the mistake is still cheap.

That qualification matters.

The cost of creating a weak prototype may be collapsing. The cost of allowing that prototype to acquire a team, customers, dependencies, operating complexity and capital is not. AI does not necessarily make wrong decisions more expensive; it gives organisations many more inexpensive opportunities to become expensively committed to them.

The prize, then, is not simply more innovation. It is **cheaper wrongness**: exposing weak propositions earlier, while they can still disappear without drama, and concentrating attention and capital behind the few things reality continues to validate.

For much of the industrial era, the company could be understood as a machine for coordinating scarce expertise and capital so that more could be produced. An AI-enabled company faces an additional problem: it may be able to generate possibilities faster than it can responsibly pursue them.

That makes the organisation resemble an editor as much as a factory.

The analogy matters because an editor does not maximise output. An editor preserves what strengthens the whole and removes what does not. For a founder, CEO or investor, the objects being edited are not sentences but commitments—products, markets, initiatives, systems, people, management attention and capital.

Follow the Scarcity

This is where AI stops being merely a technology question and becomes a question of business building and capital allocation.

A founder can explore more propositions before making a major commitment. A CEO can automate more processes and initiate more experiments. An investor looking across a company can identify dozens of plausible interventions in pricing, sales, service, product and operations.

The abundance of possibilities is valuable precisely because not all of them need to survive.

The mistake is to confuse the ability to attempt more things with an obligation to pursue more things. Every surviving initiative eventually competes for resources that AI does not make infinite: customer attention, management capacity, organisational coherence and capital.

That is the deeper inversion. For much of modern business history, intelligence itself could confer advantage because useful expertise was comparatively scarce. AI does not make intelligence irrelevant, nor does it make every form of expertise abundant. It does make some valuable cognitive capabilities dramatically easier to access, and as that happens, advantage migrates towards the constraints that remain.

Real-world evidence. Distribution. Context. Permission. Judgement. The willingness to stop.

These are not fashionable substitutes for technology. They are the things that determine whether technology becomes economically consequential.

Companies have spent decades learning how to overcome constraints on what they can produce, and much of corporate life still rewards addition: another feature, product,

market, initiative or process. AI relaxes some of those constraints without relaxing the limits on what customers will buy, organisations can absorb or investors should fund.

The defining challenge therefore changes. It is no longer simply how to make more possible. It is how to keep possibility cheap until reality has earned the right to make it expensive.

That may be the most important discipline of the AI economy.

Not predicting every model advance. Not filling companies with prototypes. Not adding AI to every product because the technology makes it possible.

Learning faster while commitment remains reversible; recognising when evidence has become strong enough to commit; and having the judgement to abandon what has not earned that privilege.

AI can help us know more, analyse more and make more. None of those things, by itself, can make customers care, businesses work or investments compound.

The technology is making possibility abundant.

Follow what remains scarce.

Roy Sharples

August 2026

Roy Sharples is the founder of Lucan Ventures, where he builds and reinvents businesses through product innovation, go-to-market and operating model design.

lucanventures.com | hello@lucanventures.com

© 2026 Lucan Ventures. All rights reserved.