

A FOUNDING VISION · 2026

THE CYBORG

Entrepreneurship

LAB

A Vision for the Future of Human-AI Value Creation

DAVID M. TOWNSEND

Digges Professor & Schulze Distinguished Professor of Entrepreneurship

Pamplin College of Business, Virginia Tech

Field Editor, *Journal of Business Venturing* · Editor-in-Chief, EIX

FOUNDING STRATEGY DOCUMENT

EXECUTIVE SUMMARY

Artificial intelligence is making high-quality thinking one of the cheapest inputs in the economy. Nearly every institution responding to this transition asks the same question – what will AI be able to do? The Cyborg Entrepreneurship Lab is founded on the opposite question: **what becomes scarce when intelligence becomes abundant?** Every prior economic transition teaches that when an input floods the economy, value migrates to whatever the flood cannot reach. Locating where value moves when thinking is no longer the bottleneck is the defining question of the next economy.

THE CONTRARIAN PREMISE. The common assumption is that more intelligence means more foresight. The Lab's research points the other way: in the domains that matter most – markets, industries, institutions – abundant machine intelligence will make the future *less* knowable, including to the machines themselves. Economic predictions are consumed by the act of using them: when millions of capable agents act on their forecasts at once, each action remakes the ground on which every other forecast stands. Quantitative finance has lived this for decades; AI will generalize it across the economy. Firms betting that AI delivers a calculable future are mispositioned for the world that is coming – faster competition, shorter-lived advantages, deeper uncertainty, and a rising premium on judgment, trust, commitment, and the capacity to act when no model can say what comes next.

WHY ENTREPRENEURSHIP. Those scarcities are precisely what entrepreneurship research has studied for a century: how people create opportunities, build legitimacy for things that do not yet exist, and exercise judgment where calculation gives out. In the intelligence age, entrepreneurship moves from business-school subfield to the discipline best equipped to understand the AI economy.

A NEW UNIT OF ANALYSIS. The replacement-versus-augmentation debate assumes humans and machines are separable performers. Increasingly, they are not. The emerging engine of value creation is the *entrepreneurial ensemble* – an integrated system of people and intelligent agents whose capabilities are properties of the configuration itself. The entity creating economic value is a human-machine composite, a cyborg in the precise sense, and it must be studied as one object: hence the Lab's name. Understanding ensembles may matter to the twenty-first century the way understanding organizations mattered to the twentieth. The questions are scattered today across fields that rarely meet; the Lab intends to define the field that gathers them.

THE RESEARCH PROGRAM. Five interlocking programs: entrepreneurial judgment in human-AI ensembles; computational entrepreneurship (large-scale simulations and synthetic economies that map possible futures rather than predict them); human-AI venture creation studied in the field; agentic organizations, including early evidence that the most aggressive AI adopters can grow more fragile even as they grow

more efficient; and entrepreneurial futures – what can be modeled, what remains unknowable, and how to act in between. The simulation architectures are built, the methods are published and peer-reviewed, and the agenda is already producing results.

THE EDUCATIONAL MISSION. When knowledge is abundant, education must teach the bottleneck – the judgment to evaluate what machines produce and the skill to design and lead human–AI systems. Graduates leave as architects of ensembles, with capabilities that appreciate as machine intelligence advances.

THE MOMENT. New fields are founded when a new kind of economic actor appears and existing disciplines cannot fully see it. The corporation occasioned organization science; the ensemble is such an actor now. The institutions that build the intellectual foundations first will define the questions a generation asks. The Lab exists to understand how humans and intelligent machines create more value together than either could alone – and to prepare people and institutions to thrive in the more uncertain, more entrepreneurial world that abundant intelligence is already creating.

THE FOUNDING QUESTION

There are moments in history when the basic assumptions underlying economic life begin to change. The Industrial Revolution transformed societies built on human and animal labor into societies powered by machines. The Information Age transformed economies organized around physical production into economies organized around knowledge. Each transition did more than change how work was done. Each one changed what was scarce, and therefore what was valuable, who prospered, and what institutions the world needed.

A third transition is now underway. Artificial intelligence is making cognitive work – analysis, research, writing, coding, design, forecasting, strategy – abundant, fast, and cheap, at a pace no previous technology has approached. Machine intelligence that once cost a salaried expert now costs pennies and runs around the clock. Within the planning horizon of nearly every university, firm, and government, high-quality thinking will be among the cheapest inputs in the economy.

Nearly every institution responding to this transition is asking the same question: what will artificial intelligence be able to do? It is the question behind the benchmark races, the capability forecasts, and the founding documents of most new AI research centers. It is also the wrong question – or at least, the question least likely to produce durable understanding, because its answer changes every few months and is largely in the hands of a small number of technology companies.

The Cyborg Entrepreneurship Lab is founded on the opposite question: **what becomes scarce when intelligence becomes abundant?**

This inversion sounds simple. It is not. It relocates the object of study from the machine's capabilities, which no university controls and no forecast can pin down, to the structure of the economic and social world the machine acts within – a world of slow-moving institutions, hard-won trust, physical constraints, human commitments, and futures that no amount of computation can fully foresee. That world is something scholars actually understand, and it is where the economic meaning of artificial intelligence will be decided. Every prior abundance teaches the same lesson. When a once-scarce input floods the economy, value does not disappear; it migrates to whatever the flood cannot reach. Mechanical power made physical strength cheap and made organization, capital, and engineering judgment dear. Information technology made data cheap and made attention, curation, and meaning dear. Abundant intelligence will be no different. The frontier question of the next economy is not what machines will think, but where value moves when thinking is no longer the bottleneck.

The Lab exists to answer that question, and to educate the people who will act on the answer.

I

THE CONTRARIAN PREMISE

More intelligence, less certainty

Beneath the founding question sits a thesis that sets this Lab apart from the capability-focused institutes now proliferating.

The most common assumption about advanced AI, shared by enthusiasts and skeptics alike, is that more intelligence means more foresight – that as machines grow more capable, the world becomes more predictable, more optimizable, more under control. Our research points to the opposite conclusion. In the domains that matter most for economic life, **the most powerful intelligence ever built will make the future less knowable, not more – including to the machines themselves.**

The reason is structural, not a limitation of today's models that the next generation will engineer away. A weather forecast does not change the weather. But an economic forecast good enough to act on changes the economy it described. When one investor finds a profitable pattern in the market, the pattern survives. When thousands of investors armed with the same analytical power find it, their trades erase it. When every firm can see the same opportunity, the opportunity is transformed by the rush toward it. Markets, industries, and institutions are reflexive: they react to what is believed about them. In reflexive domains, knowledge of the future is a depletable resource, consumed by the act of using it. The more powerful and widespread the intelligence doing the using, the faster it is consumed.

In miniature, this has already happened. Quantitative finance is the most computationally sophisticated domain in economic life, and its defining experience is that profitable patterns decay faster the more capital and computing power pursue them. The smartest money in the world has spent forty years discovering that prediction is self-consuming. Policymakers know the same lesson by another name. Measures that become targets stop measuring, and economic relationships break down once governments act on them. These episodes are previews of the dynamic that abundant machine intelligence will generalize across the entire economy.

Artificial intelligence does not suspend this logic. It accelerates it, massively. A world of abundant machine intelligence is a world in which millions of capable agents act on their predictions simultaneously, each action remaking the ground on which every other prediction stands. Forecasts will be sharper than ever and

expire faster than ever. The horizon of reliable prediction in competitive domains will contract even as predictive technology improves. This is a claim about the structure of the game, not a point forecast – closer to predicting that no one can corner a market than to predicting prices. Here lies the central paradox of the intelligence age. Capability and certainty do not rise together; past a threshold, they pull apart.



The implications reach every stakeholder a university serves. Firms betting that AI will deliver a calculable, optimizable future are mispositioned for the world that is actually coming – a world of faster competition, shorter-lived advantages, and deeper strategic uncertainty. Policymakers who expect AI to make economies more steerable will find them harder to steer. And the human capabilities that matter most will not be the ones that compete with machine prediction, but the ones that operate beyond its reach: judgment under genuine uncertainty, commitment when the spreadsheet is silent, the capacity to act wisely when no forecast can say what comes next.

Economists call uncertainty of this kind Knightian uncertainty, after Frank Knight, who distinguished a century ago between risk that can be calculated and uncertainty that cannot. The age of artificial intelligence was supposed to shrink the territory of the incalculable. We expect it to enlarge it. Defending, testing, and mapping that claim, rigorously and without sentimentality, is the Lab's central scientific commitment.

II

WHY ENTREPRENEURSHIP

The discipline of the intelligence age

If these two premises hold – abundance relocates scarcity, and intelligence amplifies uncertainty – then a striking conclusion follows about where the study of the AI economy belongs.

Consider what becomes scarce when thinking becomes cheap. Not ideas. Machine intelligence generates business plans, product concepts, and strategies in effectively unlimited supply, and the value of any input falls as its quantity explodes. What the flood cannot reach is everything required to turn a possibility into a reality. Trust, which is earned in human time and cannot be downloaded. Legitimacy, the willingness of customers, investors, regulators, and communities to accept something new. Commitment, the decision to

stake resources, reputation, and years of one's life on one venture among thousands of imaginable ones. Responsibility, which someone must bear when judgment proves wrong. And judgment itself, the capacity to act well when the analysis runs out.

These capacities form the new bottleneck of value creation. When anyone can generate a thousand venture concepts in an afternoon, the scarce act is no longer imagining the venture. It is making one real – and the economy rewards the bottleneck, not the abundance.

The scarcities of the intelligence age are precisely the subject matter of entrepreneurship research. For a century, while other fields studied optimization within known parameters, entrepreneurship scholars studied how people act when the parameters are unknown – how opportunities are created rather than found, how resources are assembled under uncertainty, how trust and legitimacy are built for things that do not yet exist, how judgment operates where calculation gives out. The field was built for a question the rest of the economy is only now being forced to ask.

This repositions entrepreneurship itself. It has long been treated as a valuable, practical, somewhat peripheral subfield of business education. In the intelligence age it becomes something closer to the paradigm discipline of economic inquiry: the field best equipped to understand where value migrates, what humans remain for, and how new value is created when intelligence is no longer the constraint. A university that recognizes this early acquires a position that capability-chasing institutions cannot replicate, because the capabilities will commoditize and the questions will not.

The Lab does not assume the entrepreneur survives this transition unchanged. Whether human judgment remains essential, where it remains essential, and where it genuinely does not – these are research questions, to be answered by evidence rather than by reassurance. An institution that only flatters the human role would produce advocacy. The Lab's value to its stakeholders depends on producing the honest map.

III

THE ENSEMBLE

A new unit of analysis

The public debate about AI and work is conducted in the language of replacement and augmentation: will the machine take the job, or help the human do it? Both framings share a faulty premise – that the human and the machine are separable performers whose contributions can be tallied on separate ledgers.

Watch how ventures are actually built today and the premise dissolves. A founder working with capable AI systems does not alternate between “human tasks” and “machine tasks.” Ideas are generated, challenged, and refined in continuous interaction. Strategies emerge from neither party alone. The reasoning is distributed across the partnership, and the capabilities of the pair are not the sum of the capabilities of the parts; they are properties of the configuration itself, for better and sometimes for worse. The same is true at the scale of the firm, where teams of humans and increasingly autonomous agents research, decide, build, and coordinate as a single operating system.

The right unit of analysis for the intelligence age is therefore not the entrepreneur, the team, or the algorithm. It is the **entrepreneurial ensemble**: an integrated system of people and intelligent agents creating value together under uncertainty, whose boundary is genuinely blurred and whose performance is emergent. This is why the Lab carries the word *cyborg* – a precise claim, with no science fiction in it, that the entity now creating economic value is a human–machine composite that must be studied as one object.

The claim has weight because the twentieth century ran this experiment once before. When the modern corporation emerged, scholars could have continued studying individual managers and treated the organization as background. Instead they made the organization itself the object of study, and organization science went on to shape a century of management practice, policy, and education. The ensemble stands where the organization stood a hundred years ago – a new kind of economic actor, spreading faster than our ability to understand it. Understanding ensembles may prove as consequential to the twenty-first century as understanding organizations was to the twentieth.

The research opening is enormous, and it is still wide open. How should an ensemble be designed – what should humans hold and what should they delegate? When does the configuration produce capabilities neither party possesses, and when does it produce failures neither would make alone? What happens to expertise, to learning, to identity, when one’s closest collaborator is a machine? How is an organization governed when some of its members are autonomous agents that never sleep, never forget, and can be copied without limit? These questions are scattered today across fields that rarely meet. The Lab intends to define the field that gathers them.

ONE QUESTION, FIVE PROGRAMS

Where value, judgment, and uncertainty move when intelligence becomes abundant

1

ENTREPRENEURIAL JUDGMENT

How judgment operates inside human–AI ensembles – which configurations sharpen it, which erode it, and what remains non-delegable.

2

COMPUTATIONAL ENTREPRENEURSHIP

Synthetic economies and large-scale simulation that map the space of possible futures rather than predict them.

3

VENTURE CREATION IN THE FIELD

Experiments and field studies across the venture lifecycle – when AI compounds entrepreneurial capability, and when it cancels it.

4

AGENTIC ORGANIZATIONS

Design, governance, and leadership of firms whose members include autonomous agents – speed without fragility.

5

ENTREPRENEURIAL FUTURES

What can be modeled, what remains unknowable, and how people and institutions act in the territory between.

*Method: a computational social science laboratory – high-performance computing, behavioral experimentation at scale, and frontier AI systems used as **instruments of research and objects of it.***

IV

THE RESEARCH PROGRAM

The Lab's agenda follows from the argument rather than preceding it: each program of work studies one face of the question of where value, judgment, and uncertainty move when intelligence becomes abundant.

ENTREPRENEURIAL JUDGMENT IN HUMAN-AI ENSEMBLES. If judgment is the binding scarcity of the intelligence age, it can no longer be studied as a property of individual minds. The Lab investigates how opportunities are evaluated, decisions made, and uncertainty navigated inside ensembles – which configurations sharpen human judgment and which quietly erode it, what experts become when expertise is abundant, and what remains genuinely non-delegable. The answers carry direct consequences for how firms deploy AI and how universities educate.

COMPUTATIONAL ENTREPRENEURSHIP: SIMULATING ECONOMIC FUTURES. Some of the deepest questions about the intelligence age cannot be answered by surveys or case studies, because the futures in question have not happened yet. Advances in high-performance computing now allow them to be studied anyway. The Lab builds large-scale simulations – synthetic economies populated by thousands of interacting human and AI agents, including simulations in which the agents are themselves frontier AI models reasoning under experimental conditions. These models are telescopes for the imagination, not crystal balls. Their purpose is not to predict the future but to map the space of possible futures, finding the thresholds where outcomes flip, the dynamics that hold across scenarios, and the early signals that distinguish one trajectory from another. Where prediction fails, rigorous cartography of the possible is the form that useful knowledge takes. The Lab's research team has already built and published with simulation architectures of exactly this kind, spanning venture creation, competitive dynamics, and organizational decision-making.

HUMAN-AI VENTURE CREATION IN THE FIELD. Simulation disciplines imagination; evidence disciplines simulation. Through experiments and field studies across the venture lifecycle, from opportunity development and customer discovery through product building, fundraising, and scaling, the Lab studies how real founders and real ensembles create value now. The early evidence already complicates the optimistic defaults: more capable AI does not uniformly produce better entrepreneurial outcomes. When every founder commands the same powerful intelligence, ventures can converge on the same opportunities and crowd away the returns; assistance that improves average decisions can suppress the outliers that drive entrepreneurial economies. Mapping when AI compounds entrepreneurial capability and when it cancels it is among the questions that matter most for firms and investors, and the Lab is built to answer it.

AGENTIC ORGANIZATIONS. Firms are beginning to employ autonomous AI agents as genuine organizational participants – researching, coordinating, transacting, deciding. These organizations differ in kind from anything management science has studied, and the Lab's early findings include a warning: organiza-

tions that adopt agents most aggressively can become more fragile even as they become more efficient, developing a surface of machine-speed rationality over a weakened human capacity to notice the novel problem no agent was trained to see. The organizations that look most transformed may be the least durably improved. Designing, governing, and leading agentic organizations well, capturing the speed without the fragility, is a defining management challenge of the coming decades, and little rigorous research exists to guide it.

ENTREPRENEURIAL FUTURES. Entrepreneurship has always been the discipline of futures that do not yet exist. The Lab's fifth program confronts the deepest version of the question: as machine intelligence advances toward and perhaps beyond human level, what about the future can be modeled, what remains irreducibly unknowable, and how should people and institutions act in the territory between? This is where the Lab's contrarian premise is put to the test at full scale – charting how far prediction reaches, how fast acting on predictions erodes them, and what strategy, policy, and entrepreneurship look like when the honest answer to “what comes next” is a map of possibilities rather than a forecast.



Across all five programs, the Lab operates as a computational social science laboratory of a new kind, combining high-performance computing, large-scale behavioral experimentation, simulation infrastructure, and frontier AI systems used both as instruments of research and as objects of it. The architecture exists, the methods are published and peer-reviewed, and the agenda is already producing results.

No single discipline can carry this agenda alone, and the Lab is built accordingly. Its questions sit at the intersection of entrepreneurship and strategy, artificial intelligence and computer science, organization theory, behavioral science, economics, and complexity science. Its working method is to put those fields in one room around shared models, shared data, and shared problems. The object of study forces the breadth: an ensemble is simultaneously a technical system, an economic actor, and a human partnership, and understanding it requires all three lenses at once.

V

EDUCATING THE ARCHITECTS OF ENSEMBLES

A research vision of this scope carries an educational obligation, because the same transition that reorganizes the economy reorganizes what education is for.

Universities have historically taught students to acquire knowledge and apply it. When knowledge itself becomes abundant, when any student can summon expert-level analysis on demand, that model quietly collapses. What cannot be summoned on demand is exactly what the Lab's research identifies as scarce: the judgment to evaluate what the machine produces, the discernment to know which of a thousand generated options deserves a life, and the skill to design and lead ensembles rather than merely use tools.

The Lab will educate for the bottleneck, not the abundance. Its students will build AI-enabled ventures as the laboratory of their own learning – leading human–AI teams, designing agentic organizations, making real decisions under real uncertainty, and studying their own collaboration with machine intelligence as they practice it. They will learn why the economics of abundance move value toward what they alone can supply, and they will graduate with capabilities that appreciate as machine intelligence advances, rather than skills that depreciate with each model release.

They will graduate, in short, not merely as entrepreneurs, but as architects of human–AI systems – capabilities every sector of the economy is already struggling to find.

The same teaching reaches beyond degree programs. Executives confronting agentic transformation, boards weighing AI strategy, policymakers designing for an economy of ensembles – all face the questions the Lab studies, years before settled answers exist. The Lab will be where they come to think clearly: a source of rigorous evidence and honest maps in a discourse dominated by hype on one side and alarm on the other.

VI

THE FOUNDING MOMENT

New fields are not founded often. They are founded when the unit of economic life changes – when a new kind of actor appears in the world and the existing disciplines, built for the old actors, cannot fully see it. The human–AI ensemble is such an actor, as the corporation and the computer once were, and the field that understands it has not yet been built.

The institutions that move first at such moments do more than produce research. They define the questions a generation asks, train the scholars who spread the field, and shape how societies understand the transition they are living through. That influence accrues only to whoever builds the intellectual foundations while the foundations are still being poured.

Fields are also built through infrastructure, not findings alone – the journals and special issues that give a research conversation standing, the conferences and workshops that assemble a community, the editorial platforms that decide what the field becomes. The Lab’s founders already hold positions of editorial and institutional leadership across the journals and outlets where the scholarly conversation about AI and entrepreneurship is taking shape, and the Lab will convene that conversation deliberately: hosting the gatherings, seeding the collaborations, and training the doctoral students who will carry the field outward. A new field needs a place where it visibly lives. The Lab intends to be that place.

The Cyborg Entrepreneurship Lab is designed for this moment. It asks the question few are asking – not what machines will do, but what becomes scarce, what becomes uncertain, and what humans become when intelligence is abundant. Its thesis, that the age of artificial intelligence will be an age of deeper uncertainty in which the scarcest resources are human ones, is one a field can be organized around. And it pairs that ambition with the computational instruments, the published track record, and the educational mission to act on it.

Its purpose is not simply to study artificial intelligence. Its purpose is to understand how humans and intelligent machines can create more value together than either could create alone – and to prepare people and institutions to thrive in the more uncertain, more entrepreneurial world that abundant intelligence is already creating.

Those questions will shape entrepreneurship, organizations, and economic life for decades.

The Cyborg Entrepreneurship Lab exists to begin answering them.

“What becomes scarce when intelligence becomes abundant?”



THE CYBORG ENTREPRENEURSHIP LAB
FOUNDING VISION · 2026
SET IN FRAUNCES & NEWSREADER