

The Conversion Problem

Why technological progress does not automatically become human capability — and what would have to be designed for it to do so.

| We are here to protect humanity, not to destroy it.

This paper argues that expanding what technology can do and expanding who can participate in it are two different engineering problems, and that only the first currently has an industry behind it. It sets out a system-design thesis, a method for testing that thesis, and the questions on which the thesis could fail.

It claims no results. The organisation described here has built nothing yet, and the paper says so wherever that matters.

ABSTRACT

A capability that exists is not a capability that is reachable, and a capability that is reachable is not one that pays. These are separated by conversions, and each conversion can fail independently. In environments that already hold records, enforceability, market access, credit and credentialed skills, those conversions happen so reliably that they become invisible — and distribution looks like a solved problem. In environments that hold few of them, the same technology produces almost nothing. The variable is neither the technology nor the person: it is the system around the technology.

We call the missing layer the **conversion layer**, and we argue it is a design problem rather than a diffusion lag. This paper states the argument, the method by which we intend to test it, and the specific conditions under which we would abandon it.

01 Two questions

The first is historical. Most of the physical and economic systems we live inside were designed under technological constraints that no longer exist. Markets, supply chains, credit, land records, labour arrangements and kinship institutions all took the shape their constraints allowed. Then technology arrived — and what followed, overwhelmingly, was *digitisation*: the existing process was preserved and given a faster interface. Paper became a PDF. A shop became a storefront. A broker became a platform. Friction fell. The underlying system did not change.

If today's technology had existed when these systems were first designed, would we have designed them the same way?

The second question faces forward, and matters more. Technological progress is compounding. Machine intelligence, automation and cheap computation will make things possible this century that are not possible now. The question is not whether that happens. It is what decides who can participate in it.

If the future becomes technologically extraordinary, what decides whether that capability reaches people who do not already hold the infrastructure required to use it?

Neither question is answered by building more technology, and that is the whole reason this paper exists.

02 The 1850 thought experiment

The historical question is easier to reason about if it is made concrete. So consider a specific and deliberately awkward version of it.

It is 1850. You have today's technology — networks, computation, machine intelligence, payments, logistics, verifiable identity. What you do not have is any of today's institutions: no agencies, no platforms, no industry has formed around any of it yet.

A family needs to arrange a marriage. The suitable families are spread across a dozen villages. Nobody can verify who anyone actually is beyond what a neighbour will vouch for. Once two families agree, months of coordination follow — negotiation, obligation, money, logistics, ceremony, and the year afterwards in which the arrangement either holds or does not.

You are designing the system. Which do you build?

- | | |
|---------|--|
| ROUTE A | Publish the matchmaker's register so families across all twelve villages can search it. |
| ROUTE B | Create profiles for eligible individuals - verified details, filters, photographs - so families find each other directly. |
| ROUTE C | Map the whole thing the two families actually have to get through, and build for the parts that fail rather than the part that is visible. |

Figure 1. Three routes from the same starting point.

Most people choose Route B. It is the reasonable answer if you begin from the process that already exists, and it is worth noticing what it produces: Route B is the matrimonial website, invented roughly a century and a half early. Discovery gets solved, and becomes very good. Verification, family coordination, negotiation, execution and the year afterwards all stay outside the system — which is where they still are.

Route A is the same error in a plainer form: the register is now searchable and the matchmaker is now faster, and nothing else has moved.

We are not claiming Route C is obviously right. We are claiming that almost nobody arrives at it by reasoning forward from the existing process, and that this pattern is not confined to marriage. Run the same experiment on agriculture, commerce, manufacturing, education, logistics, finance or the siting of settlements, and the shape repeats: the visible step gets an interface, and the system underneath is left as it was.

03 The legacy trap, in two forms

Digitising the past is the familiar version of this error, and it is well-documented enough to feel safely historical. It is not. The same mistake is available with the future, and that second form is the one worth guarding against now.




Figure 2. Both forms leave the system unchanged.

Imagine it is 2045. Machine intelligence is far more capable than today. Physical automation works. Someone looks at a fragmented economy and concludes it is behind, so the advanced technology should be deployed into it. If the system around that technology has not been redesigned — the records, the markets, the enforceability, the capital, the local interfaces — then the deployment reproduces the existing outcome with much better equipment.

The legacy trap is not only digitising the past. It is also deploying the future into an old architecture.

A technology can be radically new while the system it lands in stays old. That is why we treat the surrounding system, rather than the technology, as the design object.

04 Progress, access, capability, opportunity

This is the central analytical claim of the paper. Four things are routinely discussed as though they were one thing. They are separated by conversions, and each conversion has its own failure mode.




Figure 3. Four layers, three conversions, each able to fail alone.

LAYER	WHAT EXISTS	WHAT IT IS NOT EVIDENCE OF
Progress	A capability exists somewhere	That anyone outside the frontier holds it
Access	A device, a connection, an available model	That it can be used to do anything economically
Capability	Competent action is now possible	That the action earns anything
Opportunity	The action produces income, assets, choice	—

The error this corrects is common and comfortable: *they have smartphones and internet, so they have access to AI*. That collapses four layers into one. Access is real, and in much of India it is a genuine achievement that should not be minimised. Capability conversion is the unsolved part, and it is a systems problem — not a technology problem, and not a training problem.

The multiplier

A new technological capability is multiplied by the ecosystem surrounding the person who receives it. Give the same frontier model to a credentialed professional in a technology-dense city and to a woman running a household-scale production unit in a fragmented rural economy, and you have not given them the same thing. You have given them the same input into radically different multipliers.

TECHNOLOGY-DENSE ENVIRONMENT	INFRASTRUCTURE-THIN ENVIRONMENT
Technology	The same technology
+ existing infrastructure	- infrastructure
+ existing skills	- formal skills
+ existing capital	- capital
+ institutions and markets	- institutional connectivity
+ payment and legal rails	- enforceable contracting
HIGH CONVERSION	LOW CONVERSION

Figure 4. Identical technology. Different surrounding system. Different outcome.

Technology creates possibility. The system around it decides whether that possibility becomes capability.

A dense environment hides this by already containing the conversion layer. Records, enforceability, liquidity and credit are ambient there, so a designer working inside one can mistake them for laws of nature, build on top of them without noticing, and conclude that distribution is a marketing problem.

05 Two futures

Both futures are driven by the same technological progress. They differ only in what that progress lands on.

FUTURE A - AMPLIFICATION (well funded, works)	FUTURE B - ACCESS (largely unbuilt)
Frontier technology	Frontier technology
Existing capital, infrastructure, skills, institutions	New access architecture interfaces / trust / coordination
Higher capability	Capability where there was none
More income, more capital	Economic participation
More technology	Capital formation
[compounding advantage]	[compounding participation]

Figure 5. Both loops compound. They do not compound for the same people.

Future A compounds advantage among those who already had it. Future B compounds participation among those who did not. The important point is that **Future A is not the problem**. It is how frontier capability gets built and paid for at all; without it there is nothing to distribute.

The difficulty is that Future A running alone produces an expanding frontier with a structurally narrow entrance, and nothing in Future A's own machinery builds the entrance — because the entrance is not a

technology. It is an economic and physical system.

STATED PLAINLY, TO PREVENT THE WRONG READING

This is not an argument against technological progress, against the concentration of capital and capability that produces it, or against the organisations doing that work. It is the observation that expanding what is possible and expanding who can participate in it are two different engineering problems, and that only one of them currently has an industry. We have chosen the second, and we have not built it yet.

06 Why rural, and why that is not a deficit claim

Rural is not urban-minus. Rural and urban environments differ in kind: how trust is formed and enforced, how money and obligation move, how production is organised, which institutions are present and which absent, how information travels and how late it arrives, and what physical distance does to all of it. The deficit model — *rural is behind, ship a stripped-down city product* — is itself a legacy assumption, and it produces predictable failures that then get misattributed to the users.

We work in rural environments because that is where the conversion problem is undisguised. Dense environments conceal it by already having it. Rural environments remove the ambient scaffolding, so whatever converts a capability into opportunity has to be designed explicitly — because nothing supplies it by default.

Rural is not the limitation of the thesis. Rural is the proving ground.

Two consequences follow. First, generalisation runs in a useful direction: a system built assuming density, formality and liquidity collapses when any of those is removed, while a system built without them can usually operate where they are present. Second, the scope is larger than it appears — fragmentation, weak formal records, thin markets and trust-by-relationship describe a very large share of the world's physical economy, not a regional anomaly.

One commitment about representation, because it shapes the analysis and not only the language: rural economies are complex, productive and full of economic agency. We do not use poverty as an argument, and a design that begins from pity produces a system that condescends. Systems that condescend do not get used.

07 Designing from the outcome

If the surrounding system is the design object, the order of design has to change. The conventional order begins with the organisation and ends at the user. We invert it.




Figure 6. The interface is the last decision, not the first.

This is not a stylistic preference. Beginning from an interface commits you to the existing process, because an interface has to interface with something — and the only thing available to interface with is the arrangement that already exists. That is how Route B gets chosen.

08 Relationship to frontier technology

A great deal of the world's best talent and capital is committed to expanding what is technologically possible: intelligence, robotics, automation, computation, energy, and the frontier of where humanity can physically go. That work is legitimate, difficult and necessary. Without it there is no new capability to reach anyone at all.

What it does not answer, and is not designed to answer, is what determines who can participate once the capability exists. Progress arrives on unequal ground, and the ground is a separate design problem.

DOMAIN	THE FRONTIER QUESTION	THE QUESTION WE ADD
Space	How does humanity expand beyond Earth?	If the frontier expands, what decides who can participate in it — and is that a design choice or an accident?
Robotics	How do machines perform physical work?	How does that capability become useful to producers who cannot individually own or operate it?
AI	What can machine intelligence do?	What system lets someone in a low-infrastructure environment convert it into economic opportunity?

The robotics case is the clearest, so it is worth working through. Assume machines that do physical work become extremely capable. The default diffusion path is familiar: high capital requirement, early adoption by large operators, higher productivity, more capital, more technology. Now ask how that capability reaches a smallholder who cannot buy one. The answer is almost certainly not to give every household its own machine.

How can technological capability become economically accessible without requiring every person to own the underlying technology?

That is a question about ownership, coordination, trust and economic structure — not about robotics. It is the kind of question this organisation exists to work on, and we do not consider it answered.

09 What is being designed

Described at the level this paper can honestly support. The organisation is arranged in three parts, in dependency order.

A research layer works out what a system should become before anything is built, and tests whether the answer survives contact with the physical world. **A common foundation** provides the capabilities every real-world system needs: persistent identity, trust and verification, intelligence, a structured record of physical economic activity, coordination, interfaces and payment tied to real events. Those eight layers are precisely what a technology-dense environment has ambiently and an infrastructure-thin one lacks — which is why we describe them as the conversion layer rather than as a platform.

Independent ecosystems sit on that foundation: complete real-world systems rather than applications, interoperable but separately experienced, deliberately not merged into a single super-application.

ON WHAT THIS PAPER DOES NOT DESCRIBE

The operating model, the economics, the design of local infrastructure, and the geography in which work is concentrated are not described here and will not be. That boundary is deliberate and consistent: we intend to be open about the reasoning and closed about the mechanism. The reasoning is what can be argued with, and it is the part that benefits from being public.

10 How we intend to be wrong

A thesis that cannot fail is not a thesis. So the questions below are published with the condition that would defeat each one, and the conditions are written before any evidence exists — which is the only point at which writing them is honest.

Q1 Does new capability accumulate, or is it captured?

The argument assumes a loop: capability produces income, income produces capital, capital produces further capability. That loop can leak at every hop, and value can be captured by whoever holds the interface, the market access or the credit. If participants' measured activity rises over several cycles while their retained income and assets do not, we have produced participation without accumulation — a far weaker claim than the one made here. **This is the question most likely to invalidate the thesis.**

Q2 If capability is reached without ownership, who owns the thing it is reached through?

Separating access from ownership lowers the capital barrier and makes ownership of the intermediating infrastructure decisive. The same structure can reduce dependency or manufacture it, and we do not yet know what determines which.

Q3 Is there a minimum viable set of conditions for conversion?

Infrastructure, skills, capital, trust, market access and a physical interface all appear necessary. Which are strictly necessary, and which can be substituted — by intelligence, by coordination, by someone else's presence? The cost of the whole architecture depends on the answer.

Q4 Which parts of this arrive on their own?

Some of what we propose to build may be delivered anyway by falling costs, better models and public digital infrastructure. We cannot yet distinguish durable system gaps from temporary diffusion lags, and building for the second kind is wasted effort. This question is time-sensitive: it gets more answerable every year whether we work on it or not.

Q5 Can informal trust be represented without being destroyed?

Local trust is contextual, relational and partly unspoken. Every attempt to formalise something like it risks converting a living judgement into a score that behaves differently — which is arguably what credit scoring already did. If the act of representing reliably corrupts what is represented, the minimal useful representation does not exist.

Q6 Does coordination designed for fragmentation generalise?

Our claim that this travels beyond rural India is a hypothesis. It has not been tested, and we would rather label it than assert it.

The method, stated so it can be held against us

Each assumption is published as a research object rather than a position: with its historical origin, the strongest counterarguments made properly, three tests (would we design it this way from scratch; what does present technology change; what changes at the technological boundary), alternative

architectures compared on real criteria including reversibility, a named experiment, and a version history. The falsification condition is fixed before the experiment, and results are published either way. A result against us is the most valuable kind.

One counterargument is worth reproducing here because it is ours, and because a reader is entitled to it: **we build ecosystems, so we are predisposed to conclude that the answer is an ecosystem.** The strongest rival architecture — incremental improvement to the visible layer alone — is listed as serious rather than dismissed. What would settle it is evidence about where cost and failure actually concentrate, which we do not have.

11 What this paper does not claim

Not claimed

- No revenue, users or adoption
- No partnerships or institutional endorsement
- No valuation or market position
- No achieved scale of any kind
- No completed experiment
- No published evidence — none has been collected
- No claim that poverty will be eliminated
- No claim that this is the only possible architecture

Claimed

- That the conversion problem is real and under-addressed
- That it is a design problem, not only a diffusion lag
- That rural environments expose it most clearly
- That the questions in section 10 are open
- That we would abandon the thesis on the stated conditions
- That the reasoning is worth arguing with

The ambition is large and the work is early. Both are true at the same time, and we would rather a reader leave with an accurate picture than an impressive one. Where a document like this would normally place numbers, we have stated their absence instead.

12 An invitation, and what would actually help

This paper is being sent to people whose public work bears on questions we cannot answer alone — educators, researchers, people who have built or operated systems at scale, and people who have spent their working lives inside the environments this thesis is about.

What would help most is not agreement. Section 10 states, for each open question, the condition under which we would be wrong. A challenge aimed at one of those conditions is worth considerably more to us than endorsement is. Specifically:

- An assumption we have not noticed we are making.
- A reason a falsification condition is the wrong test.

- Evidence that one of these questions is already answered, by someone, somewhere.
- A counterargument stronger than the one we stated ourselves.
- A place where the reasoning in this paper simply does not hold.

A serious challenge gets a written answer from a person. If it changes our position, the revised assumption is published with the change noted and the contribution credited, unless the contributor would rather it was not.

CONTACT

contact@future2100.world · future2100.world

The full thesis, the open questions and the assumption tests are published at future2100.world, including the two interactive thought experiments this paper describes in sections 02 and 04.

A NOTE ON HOW THIS ORGANISATION INTENDS TO OPERATE

We are not building around announcements, coverage or the appearance of momentum. The intended order is research, system design, technology, validation against reality, capability, economic effect — and attention as a consequence of those rather than a substitute for them. This paper is circulated in that spirit: as published reasoning that can be checked, not as a launch.