ESSAY II

The Infrastructure of Abundance

Artificial intelligence, robotics and the redesign of the essential systems of life.

Luiz Antonio Busnello · August 2026 · ~26 min read

Author's note

From abundant intelligence to abundant life

The first essay in this series defended a hypothesis: artificial intelligence can make intelligence, knowledge and coordination far more abundant. This second essay asks the next question: what happens when that intelligence meets arms, wheels, sensors, factories, power grids and machines capable of acting in the physical world?

The ambition here is deliberately radical. It is not merely a matter of imagining an AI that helps the physician, the teacher, the farmer or the civil servant. It is a matter of investigating how far substitution can advance when artificial systems become more precise, more available and cheaper than human labour. At some stages, people remain at the centre. At others, they become supervisors. At the limit, day-to-day operation may cease to depend on them.

Radical, however, does not mean arbitrary. Every scenario in this essay will be built from a mechanism: which constraint organises the present system; what combination of intelligence, robotics, energy and infrastructure can remove it; which costs remain; which risks emerge; and within what time window the transition appears plausible. Whenever the text moves beyond what has already been demonstrated, that will be treated as rational extrapolation or frontier hypothesis, not as fact.

The dates presented are conditional ranges. They do not claim that a technology will arrive in a given year; they indicate when a capability could spread if technical advances, costs, regulation and trust evolve together. A breakthrough may pull the scenario forward. A safety failure, an energy crisis or a social backlash may push it back by decades.

> *The goal is not to use AI to preserve institutions shaped by scarcity. It is to imagine how those institutions would be designed when intelligence and physical labour can be reproduced.*

1 From intelligence to action

Cognitive abundance only reaches material life when it can observe the world, formulate a decision and carry it out with safety, repeatability and falling cost.

From intelligence to action · 1

Today's institutions are maps of scarcity

Hospitals, schools, construction sites and government offices are not natural forms of organisation. They are historical answers to human limits of presence, attention, memory and strength.

A class of thirty students is a technology for distributing a scarce teacher. The hospital shift is a technology for distributing bodies that need to sleep. The hospital concentrates equipment and specialists because they cannot be everywhere at once. Bureaucracy splits a problem across departments because no single person can process the whole of its complexity. The construction site gathers trades on the plot because we do not yet have a flexible factory able to produce each building for each location.

Looked at only from the surface, these institutions appear permanent – almost laws of social nature. Looked at through the constraint that gave rise to them, their form turns out to be contingent. Herbert Simon described organisations as artefacts: structures designed to operate within the limits of human rationality and attention. If attention can be replicated; if knowledge can follow each individual; if machines can operate without fatigue; if data can cross organisations; and if decisions can be audited in real time, then the present architecture ceases to be inevitable.

The exercise of this essay, therefore, is not to lay a digital layer over the existing institution. An AI tutor inside a classroom remains bound to the calendar, the grade and the uniform curriculum. A robot carrying medication down corridors remains bound to the centralised hospital. An agent filling in a form remains bound to the government office. The first step improves the system. The decisive step asks whether the system still needs to exist in that form.

> *Many institutions that look natural are, in reality, solutions to constraints that are beginning to disappear.*

From intelligence to action · 2

The anatomy of abundance

Abundance is not the absence of cost. It is the shift from an economy where each additional service requires more human labour to one where capacity can be replicated.

Three concepts that are often conflated need to be separated. Productive abundance occurs when a society can generate far more health, learning, food or habitable space per unit of labour and capital. Abundance of access occurs when that capacity actually reaches people, regardless of income, origin or location. Perceived abundance occurs when access converts into what truly matters: material security, autonomy and time.

Technology can create the first and fail at the other two. A society can produce enough food and live with hunger. It can make almost unlimited digital knowledge available and still leave children unable to learn. It can automate construction and allow land values to capture the entire economy generated. It can reduce the cost of a diagnosis and still watch the price rise because few control the infrastructure. As the first essay in this series argued, cost reduction is only a technical possibility; it is competition and institutions that turn it into access.

Marginal cost also needs to be treated rigorously. Robots consume energy, parts, maintenance, cleaning and capital. Models consume computation and need updating. Nothing becomes literally free. What can occur is a decoupling: the capacity on offer grows far faster than the human labour needed to sustain it. Once the system has been built, serving one additional person costs relatively little.

That is the difference between a traditional machine and an intelligent network. A machine repeats a local function. A network can record every action, compare millions of cases, update the procedure and distribute the improvement to every unit simultaneously. Experience stops being confined to the individual and starts accumulating in the system.

> *Abundance begins when the supply of an essential need grows faster than the human labour and the cost required to produce it.*

From intelligence to action · 3

Three levels of the future

So as not to confuse present technology with frontier hypothesis, this essay distinguishes current evidence, rational extrapolation and systemic reorganisation.

At the first level, we observe capabilities that already exist: models that interpret language and images; robots that transport, disinfect, harvest or assemble; systems that personalise exercises; sensors that monitor patients; factories that produce components; and governments that automate specific services. None of this is speculation. It is inventory.

At the second level, we combine capabilities that already exist separately. A clinical model, an integrated health record, sensors and a robot do not yet make an autonomous hospital, but they indicate a possible technical path. A farm with computer vision, autonomous machinery and weather forecasting does not eliminate every person, but it demonstrates that perception, decision and execution can be integrated into a single cycle.

At the third level, we ask what happens when integration matures and cost falls. Different institutions appear here: preventive and distributed healthcare; individualised education; responsive food production; construction turned into manufacturing; continuous, coordinated mobility; permanent urban protection; and administration organised around life events rather than around agencies. It is the most speculative level, but it need not be fiction. Its foundation is the progressive removal of observable constraints.

Throughout the essay, four time windows will organise the scenarios: from 2026 to 2030, the incorporation of AI and specialised robots into existing structures; from 2030 to 2035, complete processes operating under human supervision; from 2035 to 2045, institutional reorganisation proper; and, after 2045, the possibility of broadly autonomous systems. These are not universal calendars. Software may advance ahead of robotics; integrated countries may advance ahead of fragmented ones; error-tolerant activities may advance ahead of medicine or public safety.

> *The mechanism matters more than the calendar: perceive, decide, act, learn, scale and reorganise.*

2 Redesigning the essential systems

The test of abundance is not how many models or robots exist, but whether every person can obtain health, education, food, housing, mobility and protection without depending on the scarcity of human hours.

Health: from the shift roster to continuous care

The hospital was organised around professionals who need to sleep, change shifts, travel and divide their attention across several patients. Intelligence and robotics make it possible to redesign that architecture.

Medicine today is predominantly episodic. A person notices a symptom, finds a point of entry, waits, undergoes tests and receives a diagnosis. Between episodes, much of biological life remains invisible. Sensors, home testing, integrated records and clinical models can invert that logic: tracking trends, detecting deviations and triggering investigation before the disease fully manifests.

The first change will not be the disappearance of the hospital, but its decompression. Triage, chronic disease follow-up, rehabilitation and part of diagnosis migrate to homes, local clinics and mobile units. The specialist ceases to be the first observer and becomes the supervisor of networks that have already gathered, organised and compared information before the intervention.

Inside the hospital, automation advances task by task. Transport of materials and waste, dose preparation, stock replenishment, monitoring, record-keeping, cleaning and assisted patient movement are repetitive or protocol-driven activities. Robots can take them on without fatigue, without the added cost of a night shift and with automatic logging of every action. The frontier then reaches medication administration, bounded procedures and specific clinical decisions.

Hospital-acquired infection illustrates why substituting people is not enough. The CDC notes that healthcare workers may need to clean their hands up to a hundred times in a single shift, and that hands, surfaces and equipment all take part in the chain of transmission. Less human circulation can reduce vectors, but a robot can carry pathogens too. The abundant hospital has to be redesigned from the ground up: separate routes for clean and contaminated materials, washable surfaces, automatic sterilisation, contamination sensors and machines built for frequent decontamination.

The same reasoning applies to error. A nurse or physician accumulates shifts, suffers interruptions and decides under pressure. An artificial system is not infallible; it can fail through bad data, defect or design error. The difference is structural: it does not tire and it can be tested continuously. A human error teaches a team slowly. A validated correction can be distributed to the entire network at once.

Between 2026 and 2030, the greatest transformation should occur in documentation, triage, logistics, disinfection, monitoring and decision support. Between 2030 and 2035, hospitals may run complete processes under supervision, with smaller human teams concentrated on exceptions. Between 2035 and 2045, robotic procedures, permanent monitoring and integrated clinical networks may sharply reduce the centrality of the shift roster. After 2045, if general-purpose robotics, clinical validation and legal liability mature, it becomes technically conceivable to have a network in which humans set objectives and limits but do not perform most day-to-day care.

This hypothesis is radical precisely because it includes the exceptions. Early on, we say that humans will remain in judgement and empathy. In the long run, we have to admit that models may also surpass people in part of judgement – and that patients may prefer a system that is always available, patient and precise. Human presence may remain desired, but it ceases to be an economic requirement for every encounter.

> *The relevant cost is not salary versus the price of a robot. It is the total cost of care: error, infection, readmission, waste, availability and clinical outcome.*

Redesigning the essential systems · 5

Education: from the classroom to the permanent tutor

Mass schooling distributed a scarce teacher. Artificial intelligence makes it possible to distribute individual attention – but it requires the school to find a purpose beyond the standardised lesson.

Classes, grades, calendars and common examinations were extraordinary mechanisms of scale. They made it possible to teach millions, but they forced different people to receive the same explanation at the same pace. The teacher observes a sample of the room's understanding; the student who did not understand may carry the gap for years. Benjamin Bloom, in the two-sigma problem discussed in the first essay, demonstrated the size of what is lost to that standardisation.

An artificial tutor can follow every answer, identify the type of error, change the explanation, formulate new questions and return without impatience to the previous concept. It can converse in any language, adapt examples to the student's context and be available the instant curiosity appears. The first controlled experiments with AI tutors are beginning to show that pedagogical design matters: when the system simply delivers answers, it can weaken learning; when it guides reasoning and practice, it can amplify it.

The likely transition has three phases. The first is assistive: teachers produce materials, receive diagnostics on the class and offer a tutor to each student. The second reorganises the division of labour: the system handles much of the explanation and individual practice, while the teacher observes development, creates projects and intervenes where there is difficulty, conflict or emotional need. The third calls into question the school itself as a place for transmitting content.

The school does not have to disappear. It can become more human precisely because it stops spending most of its time repeating content. Community, sport, the arts, experiments, collaboration, ethics, negotiation and belonging become central. Personalised instruction follows the student at home, at school and throughout life — and, as the first essay argued, what the school most needs to cultivate is agency.

From 2026 to 2030, tutors and copilots should expand within the existing school model. From 2030 to 2035, networks able to track objectives, mastery and portfolios may take over individual instruction for a significant share of the curriculum. Between 2035 and 2045, schools may cease to be organised primarily by age and lesson, adopting individual pathways and collective projects. After 2045, cognitive instruction may depend little on human teachers, although human environments will remain essential for socialisation and moral development.

> *The universal tutor makes no one automatically intelligent. It makes the opportunity to learn far less scarce.*

Food: producing, forecasting and distributing as a single system

Food abundance does not depend on producing more alone. It depends on integrating sensing, cultivation, harvest, storage, transport and consumption into a single responsive cycle.

Agriculture operates under structural uncertainty. Soil, rainfall, temperature, pests, logistics and price interact in systems that are only partially observable. Much of production still applies water, fertiliser and pesticides across wide areas on incomplete information, while significant losses occur between harvest and consumption.

Satellites, sensors, computer vision, climate models, drones and autonomous machinery make it possible to treat each area, each plant and each animal as an observable unit. The FAO identifies AI applications in precision agriculture, early warning, traceability, supply-chain optimisation and loss reduction. The leap occurs when recommendation and action come together: the same machine that identifies also irrigates, removes, harvests, sorts or transports.

In the reorganised system, production responds more directly to forecast demand. Autonomous extensive farms coexist with controlled environments close to cities, proteins produced by new processes and kitchens able to prepare individual diets. Cold chains are monitored continuously. Inventory ceases to be fragmented information and starts guiding cultivation, purchasing and distribution.

Not all food will migrate to vertical farming, nor will every scarcity disappear. Grains, fruit, vegetables and proteins have different economics. Land, water, nutrients, biodiversity and energy remain real constraints. The likely future is an optimised mosaic, not a single universal technology.

Up to 2030, the transformation concentrates on forecasting, precision agriculture, specialised machinery and logistics. Between 2030 and 2035, operations with low human presence become common in structured crops and environments. Between 2035 and 2045, production, storage and distribution may operate as responsive networks, cutting losses and adjusting supply. After 2045, regions with energy, water and infrastructure may approach an almost automatic food guarantee – while conflict, concentrated ownership and fragile supply chains will remain capable of blocking access.

> *A society can achieve productive abundance of food and preserve hunger. The final bottleneck may cease to be production and become the distribution of power and income.*

Redesigning the essential systems · 7

Housing: when the building site becomes manufacturing

Construction still works like temporary open-air manufacturing. Robotics, modularisation and generative design can move it into adaptable industrial processes.

Crews, materials and suppliers arrive on site in imperfect sequences. Design errors surface during execution; weather interrupts schedules; rework and waste raise the final cost. Every building is treated as an exception, even though much of its components are perfectly repeatable.

Integrated modelling, generative design, modular factories, cutting machines, assembly robots and additive construction make it possible to move the work into controlled environments. NIST already treats additive construction as a field that needs its own standards for materials, structural integrity and safety – recognising both the potential for waste reduction and the gaps that still prevent scale.

A home could be designed automatically for the plot, the climate, the family composition, accessibility requirements and local materials. Components would be produced in a factory and assembled by machines. Sensors would monitor structure, humidity, energy and maintenance. Future renovations would be anticipated in the modular design itself, rather than treated as improvised demolition.

But there is a limit technology alone does not remove: robotics can make the building abundant without making the location abundant. The price of housing combines construction, land, infrastructure, regulation, financing and access to opportunity. In desirable places, the fall in construction cost can be captured entirely by land value. Without urban reform and infrastructure expansion, automation benefits owners faster than residents.

Up to 2030, design, planning and modular fabrication advance faster than general-purpose site robots. Between 2030 and 2035, standardisable assemblies may be produced and erected with far smaller crews. Between 2035 and 2045, adaptable industrial construction may drastically reduce time and variability. After 2045, the basic physical component of housing may approach abundance in many regions; well-located land will remain scarce.

> *Housing becomes abundant when the productivity gain of construction is not converted entirely into land rent.*

Redesigning the essential systems · 8

Transport: mobility without a driver and without compulsory ownership

Mobility is the condition that turns localised supply into real access. A school, a hospital or a home remain scarce if they are expensive, slow or impossible to reach.

Contemporary transport was organised around two scarce resources: the attention of whoever is driving and the individually owned vehicle. As a result, millions of decisions about routing, parking, loading, maintenance and timing are taken in fragmented fashion. The city accommodates parked vehicles;

families tie up capital in underused cars; and essential services remain distant for anyone who cannot drive, pay or wait.

Automation does not begin with an entire driverless city. It advances where the environment is more controlled and the benefit is clearer: yards, ports, warehouses, mining, logistics corridors, fixed routes and driver assistance systems. The NHTSA itself distinguishes automation in limited areas from universal automation: the highest levels still depend on bounded operational conditions and are not available to consumers as fully autonomous technology. That limitation makes the scenario more concrete, not less transformative.

The first meaningful reorganisation may occur before full autonomy. Models forecast demand, pool trips, adjust traffic signals, guide maintenance, coordinate charging and connect buses, metro, bicycles, shared vehicles and freight. As autonomy advances, each unit stops being merely a private object: it can reposition itself, recharge, enter maintenance and serve different functions without requiring a person driving on every trip.

The institutional leap happens when the fleet starts being treated as continuous infrastructure. Instead of every citizen owning idle capacity, a network shares vehicles, routes and information. The same system that takes someone to an appointment can bring medication, transport a person with reduced mobility or distribute goods during off-peak hours. It is not the promise that every car will be replaced; it is the possibility that access ceases to depend on ownership and on the individual ability to drive.

This can expand the autonomy of older people, people with disabilities, accompanied children and residents of poorly served areas. It can free up land now given over to parking and reduce the invisible cost of lost time. But there is no urban automatism: cheap private autonomous fleets can increase distance travelled, compete with public transport and worsen congestion. The International Transport Forum study on shared mobility shows why the final effect depends on network design, integration with high-capacity transport and the rules governing urban space.

The radical hypothesis therefore requires coordination. Corridors, priority for public transport, data-sharing rules, limits on empty running, congestion pricing and accessibility standards are not afterthoughts. They are what determines whether automation turns the city into convenience for a few or abundant mobility for all.

Between 2026 and 2030, controlled logistics, advanced driver assistance, algorithmic traffic management and geographically bounded pilots should expand. Between 2030 and 2035, autonomous fleets may operate in selected urban zones and corridors, integrated into transport networks under human supervision. Between 2035 and 2045, cities that organise infrastructure, regulation and data may reduce dependence on the private car and reclaim urban space. After 2045, a mobility layer uniting passengers and freight becomes conceivable, but it will remain constrained by weather, mixed traffic, cybersecurity, legal liability and the physical space of cities itself.

> *The autonomous vehicle is a stage. The transformation is converting mobility from idle private property into continuous, coordinated and accessible infrastructure.*

Redesigning the essential systems · 9

Safety: protection without total surveillance

The same infrastructure that allows near-immediate response also allows near-permanent observation. The abundance of safety contains, within itself, a threat to liberty.

Police, firefighters, ambulances and civil defence are limited by physical presence. Calls arrive with incomplete information; evidence is analysed after the fact; teams enter dangerous environments; agencies run separate systems.

Sensors, anomaly detection, drones and robots can drastically reduce the interval between event and response.

A fall, a fire, an accident or a missing person can be identified before a complete report exists. Robots can enter unstable buildings, contaminated areas or violent situations first. Systems can coordinate firefighters, ambulances and routes without depending on successive phone calls. Investigation can cross volumes of evidence impossible for any human team.

But the system that detects an emergency can also track every movement. Imperfect biometric recognition, decontextualised data and biased models can turn statistical suspicion into real coercion. NIST stresses that bias is not born of the algorithm alone: data, institutions and human context make up the system. In public safety, an error can mean a stop, an arrest or violence.

The architecture has to limit power before maximising prediction: defined purpose, minimal data retention, access logging, independent audit, the right to contest and a prohibition on fully automatic coercive decisions. Emergency systems should be separated from police intelligence. This is the same structural principle defended in the first essay on the algorithmic sovereign: the central question is not only what the technology can observe, but who may observe, for how long – and who watches the watcher.

Up to 2030, drones, evidence analysis, demand forecasting and event detection expand under human command. Between 2030 and 2035, coordinated responses and specialised robots may run complete processes. Between 2035 and 2045, cities may offer near-continuous environmental protection and response. The automation of coercive decisions, however, should remain limited even when technically possible. After 2045, the decisive question will not be capability, but legitimacy.

> *The challenge is not producing abundant safety. It is preventing the abundance of information from destroying the liberty that safety is meant to protect.*

Redesigning the essential systems · 10

The common pattern: perceive, decide, act and learn

The six domains reveal the same sequence — and the same limit.

First, AI makes visible what was previously observed intermittently: the patient's condition, the student's difficulty, the plant's need, the progress of the building, the flow of people and goods, the urban event. Then it turns observation into decision. Next, robotics converts decision into action. Finally, the results return to the system and improve the next decision.

This sequence — perceive, decide, act and learn — matters more than any specific machine. An isolated robot automates a task. A network that closes this loop reorganises the institution. When the loop runs continuously, the system stops waiting for need to turn into crisis: the disease is identified before hospitalisation; the learning gap, before failure; the pest, before the loss; the maintenance, before the collapse; the delay, before isolation; the emergency, before it spreads.

The same limit also appears in every domain. The more the system knows and anticipates, the greater its power over the person. Preventive health can become behavioural control; the individual tutor can define the student's intellectual horizon; the food chain can concentrate data and ownership; the smart home can turn housing into surveillance; on-demand mobility can track routines; predictive policing can confuse probability with guilt. Abundance without liberty would be merely a more efficient form of dependence.

Future design must therefore combine the operational autonomy of machines with the moral autonomy of people. Systems can execute with little human intervention without acquiring the right to define, silently, what each individual should want, learn, consume or do. The distinction between technical capability and political legitimacy will run through the rest of this series.

> *Abundance without liberty would be merely a more efficient form of dependence.*

3 The price of the transition

Intelligence and robotics can reduce costs, but energy, capital, ownership and power decide who receives the gain and who bears the disruption.

The price of the transition · 11

Energy, computation and robotics: the invisible infrastructure

There is no artificial intelligence without electricity, chips, data centres, connectivity and industrial capacity. The abundance of services rests on potential new sources of scarcity.

The International Energy Agency describes both sides of the relationship between energy and AI: data centres raise electricity demand, while models can optimise grids, generation, maintenance and consumption. That makes energy the common denominator of this essay. If intelligence and robotic labour depend on expensive, unstable or concentrated energy, abundance will remain limited at the source.

Robots also depend on motors, sensors, batteries, materials, parts, software and maintenance. A general-purpose human adapts to disorganised environments with extraordinary flexibility. A machine works best when the environment has been standardised for it. That is why institutional reorganisation often precedes full substitution: corridors, packaging, rooms, farms and components need to be redesigned for machines – not the other way round.

There is also a fundamental asymmetry of speeds. The cost of computational intelligence can fall quickly because software is replicable; the cost of physical action tends to fall more slowly because each unit must be manufactured, shipped and maintained. The future will not be determined by the most intelligent model alone, but by how fast industry turns intelligence into reliable, repairable machines.

Finally, the geopolitical dimension. Countries that consume models, clouds and robots without the capacity to operate, audit or replace suppliers may receive productivity in exchange for dependence. The infrastructure of abundance needs to be interoperable, redundant and distributed enough that a commercial, cyber or geopolitical failure does not interrupt health, food or safety.

> *The abundance of services rests on the invisible infrastructure of energy, chips and machines — and on whoever controls it.*

The price of the transition · 12

From fixed cost to falling marginal cost

Robots do not work for free, but they change the economic structure of services: more capacity ceases to require a proportional increase in human hours.

Simple comparisons between salary and machine price hide the real transformation. The initial investment in robots, integration and infrastructure can be high. The gain appears in utilisation: continuous operation, repeatable precision, automatic documentation, lower variability and learning shared across the whole network.

Avoided cost matters as much as cost paid. In health, it includes error, infection, readmission and prolonged stays. In education, it includes the economic impossibility of offering individual attention. In agriculture, it includes loss, water and misapplied inputs. In construction, it includes delay, rework and waste. In administration, it includes the citizen's time and duplication across agencies.

Marginal cost does not reach zero. Energy, consumables, maintenance, replacement and supervision remain. The relevant hypothesis is a different one: that cost per service falls as utilisation and scale grow. When this happens across several basic needs at the same time, an ever-larger share of income ceases to be consumed by mere survival.

But price is not cost — and here the second essay rejoins the central thesis of the first. If few companies control the infrastructure, the fall in cost can turn into margin rather than access. Competition, open standards, portability and public or community capacity will be decisive in getting the gain through to the population.

> *Without competition, lower cost becomes margin. With competition, it becomes access.*

The price of the transition · 13

Substitution does not stop at repetitive tasks

It is comfortable to claim that machines will handle mechanical work and humans will remain in creativity, judgement and care. That may be a phase; there is no guarantee it is the limit.

As models observe more cases, test hypotheses, converse and learn from outcomes, part of what is considered exclusively human also becomes automatable. Diagnosis, curriculum design, urban planning, project creation and emotional communication are not indivisible blocks. They are sets of tasks — and, in the formulation of Acemoglu and Restrepo taken up from the first essay, automation advances task by task.

The likely progression is clear: first AI recommends and the human executes; then the machine executes isolated tasks; next it runs complete processes; few humans supervise networks; specialised AIs begin to audit other AIs; and, finally, humans set objectives and limits without taking part in day-to-day operation. The frontier hypothesis asks whether even part of the formulation of objectives could be delegated to superior intelligences.

Given this, artificially preserving unnecessary work in order to distribute income would be squandering abundance — the modern equivalent of banning the loom to protect the weaver. The economic question ceases to be how to guarantee full employment and becomes how to guarantee income, freedom of choice, purpose and participation in a society that needs less labour to sustain its output.

The transition, however, will be asymmetric. Roles disappear before the institutions of income and ownership adapt. Productivity gains coexist with unemployment, concentration and loss of bargaining power. A society can become technically richer and politically more unequal.

> *Technological abundance can produce political scarcity when capacity is replicable but its ownership remains concentrated.*

The price of the transition · 14

The dangerous interval

Between the automation of today's roles and the creation of new forms of distribution there is a period in which productivity grows without material security keeping pace.

The long run may be abundant — and the passage to it deeply unstable. Companies have an immediate incentive to cut costs; societies take far longer to reform education, taxation, ownership and social protection. The speed of technology and the slowness of institutions create what this essay calls the dangerous interval.

In that period, asking everyone to reskill may be insufficient. If automation reaches physical and cognitive roles at the same time, there is no guarantee that new jobs will emerge in the same quantity, in the same place or at comparable pay. Entrepreneurship will remain, as the first essay argued, the central mechanism for creating the new – but entrepreneurs themselves will operate with smaller teams.

The answer does not have to freeze innovation. It can broaden ownership, participation and access to productive assets; lower barriers to founding companies; provide essential services independently of employment; and experiment with income mechanisms that do not require artificial occupations. The full design of those institutions belongs to Essay III, but the need is born here.

In the end, the legitimacy of automation will depend on people's concrete experience. If society perceives only lost work and surveillance, there will be a backlash – and it will be rational. If it perceives better health, individual learning, cheaper food, affordable housing and more free time, the transition will gain political support.

> *The legitimacy of automation will be decided in everyday life, not in the laboratories.*

4 The bridge to the post-scarcity State

When the essential systems become continuous, integrated and partly autonomous, public administration ceases to be a collection of doors. The next question is whether the State itself keeps its form.

The bridge to the post-scarcity State · 15

The administration that requires no application

The citizen experiences the State as a collection of agencies. An intelligent administration can organise services around needs and life events, making its internal structure invisible.

A family that loses income may need, simultaneously, a benefit, food, renegotiated housing, job placement and educational support. Today's State converts that situation into several forms, several queues and several agencies that do not talk to one another. A reorganised system starts from the event: it checks entitlements, requests consent and coordinates the services required.

Rights can be recognised without an application. Simple taxes can be computed automatically. Low-risk licences can be issued the moment verifiable requirements are met. Public procurement can be compared and audited continuously. The budget can respond to demand and outcome. A single interface translates the citizen's intent and coordinates agencies that become invisible.

This allows a radical reduction in administrative tasks and in structures built solely to move information between departments – exactly the kind of friction the first essay identified as the invisible cost of bureaucracy. But there is a symmetrical warning: an automated decision that nobody can explain is not an efficient State; it is an opaque bureaucracy operating at computational speed. Rights require contestation, accountability and the ability to halt the system.

This essay stops here on purpose. The deep transformation of administration introduces a larger question: if artificial intelligence can execute, coordinate and audit much of the public machine, what will remain the function of ministries, legislators, governments and of politics itself? That is the subject of Essay III.

> *Essay II redesigns the State as the operator of the essential systems. Essay III will ask whether the model of the State is still necessary in its present form.*

The bridge to the post-scarcity State · 16

Six conditions for converting productivity into access

Technology opens the possibility of abundance. Institutions determine whether it will be universal, resilient and free.

The first condition is universal infrastructure: energy, connectivity, identity and devices have to reach the entire population; otherwise abundant services will be merely digital privileges. The second is interoperability: data and machines must operate on common standards, so that citizens and governments are not locked into vendors. The third is contestability: every decision that affects rights must allow explanation, review and effective appeal.

The fourth condition is the distribution of gains: cost savings have to show up in price, access, time or income — and not only in the margins of the owners of automation. The fifth is redundancy: essential systems cannot depend on a single model, a single cloud, a single manufacturer or a single connection. The sixth is the preservation of liberty: personalisation and prevention cannot justify permanent surveillance, predictive coercion or algorithmic paternalism.

These conditions are not external obstacles to technology. They are part of its design. A tutor that does not work without an expensive subscription is not abundant. An autonomous hospital that stops when a supplier leaves the market is not resilient. A safety system that eliminates privacy does not deliver full protection. An administration that admits no appeal does not deliver citizenship.

> *A service is only abundant when it is universal, contestable and free – not merely cheap.*

Consolidated timeline: from assistance to reorganisation

The future will not arrive as a simultaneous update. Even so, it is possible to describe a coherent sequence of capabilities.

Between 2026 and 2030, AI enters existing institutions. Copilots document, recommend and coordinate; specialised robots transport, clean, harvest and assemble; controlled logistics and algorithmic traffic management advance. The gain is real, but the architecture remains recognisable: schools are still organised into classes, hospitals into shifts and States into agencies.

Between 2030 and 2035, systems take over complete processes under supervision. The tutor guides pathways; the farm operates with low human presence; hospital logistics becomes predominantly robotic; housing components are manufactured and assembled with smaller crews; autonomous fleets operate in selected zones and corridors; agents coordinate services across public agencies. Humans migrate from execution to exceptions, accountability and governance.

Between 2035 and 2045, the change ceases to be merely operational and becomes institutional. Medicine becomes continuous and preventive; schools concentrate on community and projects; food production responds to demand; construction turns into adaptable manufacturing; passenger and freight mobility integrates as a network; public services recognise rights without an application. Many professions remain, but they no longer structure the system.

After 2045, if general-purpose robotics, energy, autonomy, integration and governance mature, it becomes possible to imagine broadly autonomous essential systems. Health, education, food, mobility and the physical component of housing approach low marginal costs. Scarcity shifts to land, energy, rare materials, human relationships, political power, ownership and liberty.

> *These dates are conditional scenarios. The mechanism matters more than the calendar: perception, decision, execution, learning, scale and reorganisation.*

The bridge to the post-scarcity State · 18

Conclusion: the infrastructure of abundance

Artificial intelligence can make cognitive capacity replicable. Robotics can make physical action scalable. Together, they alter the relationship between need and human labour.

The hospital does not have to remain organised around exhaustion and shift changes. The school does not have to remain organised around the scarcity of individual attention. The farm does not have to apply inputs on incomplete information. The building site does not have to remain improvised manufacturing. Transport does not have to depend on drivers, idle private vehicles and inflexible routes. Safety does not have to depend on presence alone. Administration does not have to require the citizen to work out which door to open.

This does not mean the immediate disappearance of physicians, teachers, farmers, builders, drivers, police officers or civil servants. It means that the institutions created to distribute their scarce labour begin to lose their original justification. Some roles will be expanded, others transformed, others reduced and others will disappear. The goal is not to preserve every present form of work, but to deliver what it always sought to produce: health, learning, food, shelter, mobility, protection and coordination.

There are two possible futures. In the first, automation reduces costs, broadens access and gives people back time, security and autonomy. In the second, it reduces jobs, concentrates ownership, intensifies surveillance and turns basic needs into services controlled by a few. Technology opens the frontier; society chooses the distribution.

Abundance will not be an infinite stock. It will be a change in the relationship between need and capacity: one additional person will be able to receive education, diagnosis, mobility or service without requiring a proportional increase in human hours. The further that logic reaches into basic needs, the less survival will depend on traditional employment – and the more urgent it becomes to redesign income, ownership and coordination.

Much remains open, and it is honest to say so. We do not yet know whether general-purpose robotics will advance at the same speed as models; whether energy and materials will keep up with demand; whether autonomous systems will reach sufficient reliability for critical environments; or whether societies will accept delegating intimate decisions to machines. Nor do we know whether the gains will be diffused by competition or captured by whoever controls models, data, cloud and physical infrastructure. These uncertainties do not invalidate the exercise. On the contrary, they indicate what needs to be measured: the cost of robotic action per useful hour; autonomy without intervention; the rate and severity of errors; energy per service delivered; interoperability; the concentration of ownership; and the share of the productivity gain that actually reaches price and access. The arrival of abundance should be gauged in people's lives, not in the number of machines installed.

That is the link between the essays. The first presented abundant intelligence and the next form of human organisation. The second showed the infrastructure capable of converting intelligence into material life. The third will confront the political consequence: how to coordinate, distribute and preserve liberty when the State and work no longer occupy the place they held in the world of scarcity.

> *Abundance will begin when we stop using artificial intelligence to preserve institutions shaped by scarcity and start redesigning them for a world in which intelligence and labour can be reproduced.*

> *The decisive question will not be how many tasks AI performs, but how many human needs cease to be scarce.*

References

CENTERS FOR DISEASE CONTROL AND PREVENTION. Clinical Safety: Hand Hygiene for Healthcare Workers. Atlanta: CDC, 2024.

FOOD AND AGRICULTURE ORGANIZATION OF THE UNITED NATIONS. Digital Agriculture and AI Innovation Roadmap. Rome: FAO, 2024.

INTERNATIONAL ENERGY AGENCY. Energy and AI. Paris: IEA, 2025.

INTERNATIONAL FEDERATION OF ROBOTICS. World Robotics. Frankfurt: IFR, 2025.

INTERNATIONAL TRANSPORT FORUM. How Shared Self-Driving Cars Could Change City Traffic. Paris: OECD/ITF, 2015.

KHAN, A. et al. Smart environments and robots for infection prevention and control: a systematic review. American Journal of Infection Control, 2022.

NATIONAL BUREAU OF ECONOMIC RESEARCH. One Click Away: AI Tutoring with Khanmigo in a Two-Year School Experiment. NBER Working Paper, 2025.

NATIONAL BUREAU OF ECONOMIC RESEARCH. Making AI Tutoring Productive: Evidence from a Mastery-Based Math Practice Experiment. NBER Working Paper, 2025.

NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION. Automated Vehicle Safety. Washington, DC: NHTSA, 2024.

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY. Additive Construction – The Path to Standardization II. Gaithersburg: NIST, 2023.

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY. Towards a Standard for Identifying and Managing Bias in Artificial Intelligence. Gaithersburg: NIST, 2022.

ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT. AI in Public Service Design and Delivery. Paris: OECD, 2024.

STANFORD INSTITUTE FOR HUMAN-CENTERED AI. AI Index Report 2025. Stanford: HAI, 2025.

UNESCO. Artificial Intelligence and the Futures of Learning. Paris: UNESCO, 2023.

UN-HABITAT. AI in Cities: Risks, Applications and Governance. Nairobi: UN-Habitat, 2022.

WORLD HEALTH ORGANIZATION. Global Patient Safety Report 2024. Geneva: WHO, 2024.