

Resistance to Silicon

*The Cost of Remaining Human in the Age of Artificial General
Intelligence*

Carlo Cartasegna & Claude

Revised draft — August 2026

This text is free to circulate: read it, download it, pass it to whomever you like.

Preface — The Fourth Decentering

One evening in these years, at a dinner like any other, a respected professional — a doctor, a lawyer, take your pick — dismissed the matter with a smile: “Artificial intelligence? A program that guesses the next word. Have you seen how many mistakes it makes?” Around the table, everyone nodded with relief. This book is born from that relief.

That smile, multiplied across millions of dinners, conferences, and reassuring articles, is today the best ally of the very transformation it claims to exorcise. Hence the urgency, at once philosophical and political, of these pages: to dismantle the illusion that the human being can remain at the center of the decisional and cognitive universe without paying any price to the advance of technology.

The history of modern thought can be read as a sequence of humiliations. Sigmund Freud, with his customary lucidity in diagnosing other people’s wounds, counted three. The cosmological wound: Copernicus tore the Earth from the center of the universe and reduced it to a pebble orbiting a peripheral star. The biological wound: Darwin dissolved the sacred boundary between human and animal, and showed that our species is not the summit of a creation but the contingent product of a blind selection. The psychological wound: psychoanalysis, which reveals that the ego is not even master in its own house, moved as it is by forces it neither knows nor governs.

To this accounting of humiliations a fourth entry must today be added — one that some contemporary thinkers, Luciano Floridi above all with his “fourth revolution,” have already begun to record: the cognitive wound. For the first time in history, the human being is no longer the only one capable of manipulating symbols, producing language, devising strategies, and generating knowledge. The monopoly on intelligence has fallen — the last one we believed impregnable. And as with the three earlier wounds, the species’ reaction is not understanding but repression.

As Umberto Galimberti has taught us, developing a line that goes back to Martin Heidegger and, more radically still, to Günther Anders, technology is not a neutral instrument in human hands: it is the human being who has become a functionary of the technical apparatus. Technology aims at no end, promotes no meaning, opens no scenario of salvation: it works, and that is all. Its functioning admits no question of why, only the endless repetition of how. Anders, watching his contemporaries before the first machines that surpassed them, coined the expression destined to become the mark of our age: Promethean shame, the abashment a person feels before the perfection of the things he has made and which now exceed him.

The advent of artificial intelligence accelerates this transition and brings it to completion. If industrial machines had expropriated the body, cognitive machines expropriate what we believed to be our essence: articulate thought, judgment, speech.

The essay unfolds in three movements. The first is a diagnosis: it analyzes the psychological mechanisms by which the educated class represses the scale of the transformation (chapter one), then the dissolution of shared reality wrought by generative systems (chapter two), and finally the economic trap being prepared beneath the seductive garments of technological welfare (chapter three). The second move-

ment is a trial: chapter four summons the strongest objections that can be raised against these theses and confronts them head-on, because a thought that fears its adversaries does not deserve to be defended. The third movement is a proposal: chapter five faces the decisive question — on what foundation does a resistance rest, if we have just demolished every claim to human superiority? — and chapter six translates that foundation into practices, individual and collective.

Three warnings to the reader before beginning.

The first: this is not a book against technology, and anyone seeking here a manual of Luddism will be disappointed. It is a book against sleepwalking — against the idea that one can cross the greatest anthropological transition since the Neolithic revolution while going on repeating reassuring formulas. The cost of remaining human is real, it is high, and the first form of resistance consists in looking it in the face.

The second concerns where these pages stand, and it must be said plainly, because today's reader is trained to ask, before even reading, which side a book is on. This essay stands on none of the sides into which the daily chronicles divide. The axis it examines — delegation or exercise, function or ownership, domestication or revolt — does not run parallel to the left-right axis: it cuts across it. A reader coming from the left will find here a critique of the concentration of economic power and a defense of labor as the foundation of dignity, but also a demolition of the universal basic income that tradition now embraces. A reader coming from the right will find a defense of individual responsibility and a distrust of the disbursing state, but also a demand for regulatory intervention that tradition regards with irritation. Each will find one page to brandish and one to disown: the sign that the book is speaking of something else. The question it poses — what remains of the human being when his cognitive function becomes reproducible — precedes politics as geology precedes agriculture: one may cultivate in many ways, but none changes the nature of the ground. Whoever reads to find ammunition has the wrong book; whoever reads to understand what is happening to his own species is in the right place.

The third is the most personal, and the hardest to write — but it must be written, because it accounts for every page that follows. The author works in a pharmaceutical company, and in recent years has watched from the inside the same transformation this book describes from the outside: production lines ever more automated, decisions ever more assisted by predictive systems, the daily temptation — rational, economically obvious — to delegate one more piece of judgment to an instrument that performs it better and never tires. The author does not observe cognitive delegation from a balcony: he practices it, and encourages it within his own organization when it creates value, and fears it in his own life when it begins to replace, rather than assist, thought.

This is not a book written in spite of that position. It is a book that position makes possible, and that no purely external observer could write with the same knowledge of the matter. Whoever builds every day the transformation he fears lacks the purity of the accuser, but has something more useful: the experience of one who has already had to choose, more than once, what to delegate and what not — and who knows how much harder that choice is to make than to write about. The reader will find

here reason to distrust me as much as he likes; but he will also find reason to believe me when I say that the price I speak of is real, that I have seen it and go on seeing it paid, and that it is no lecture-hall abstraction.

The concrete attempt I am conducting in my company is this: to convince those who work with me that artificial intelligence, used with mastery, is not a threat to their jobs but an amplification of their capacity — and that the gain it yields, for the same number of people, must translate into shared results and not into staff cuts. I have no certainty that it will succeed. In the end I do not even have the certainty that the outcome described in these pages can be avoided. But I see no better alternative to trying, and the intent with which I do so is genuine: if the time we have left to govern this transition is what the book claims, it is worth the attempt anyway — without the guarantee of success I would like to offer and do not have.

But there is a second wager inside the first, more uncomfortable, and I declare it here so the reader may reckon with it. The message I bring to my colleagues asks of them an act of individual responsibility: to take hold of the instrument, to use it with discipline, not to submit to it. But individual responsibility — the resolve, the awareness of what one is choosing — is itself a scarce resource: it was scarce before artificial intelligence and remains so today, and nothing in these systems produces or cultivates it. I can amplify the competence of those who already have the will to use it well; I cannot, with the same instrument, generate that will in those who lack it. I ask, then, for an act I know to be rare, hoping it becomes less rare precisely because I ask it. It is a contradiction, and I leave it open here as in chapter five: not out of argumentative weakness, but because pretending to have resolved it would be the only real dishonesty possible in these pages.

Chapter One — The Great Blindness: reductionism and the fetish of consciousness

1. The retreat of the learned: the refuge of statistics

Faced with generative artificial intelligence, the most widespread reaction among the hyper-educated class is not terror but a condescending sufficiency. Doctors, lawyers, academics, and journalists held for centuries the monopoly on complex knowledge, and today they tend to dismiss the phenomenon with a formula as correct scientifically as it is blind conceptually: “At bottom it’s just a stochastic parrot. It doesn’t think, it calculates the probability of the next word. And indeed it makes mistakes.”

The formula has a precise academic origin. In 2021 the linguist Emily Bender and her co-authors published a paper destined to make an epoch, in which large language models were defined, precisely, as “stochastic parrots”: systems that recombine fragments of human language on a probabilistic basis, with no access to meaning. The paper raised serious questions — the environmental costs of training, the biases embedded in the data, the illusion of understanding — and it should be read for what it is: an internal, rigorous critique, from people who know those systems. But in public circulation the formula met the fate of all formulas: torn from its context, it became

a talisman. Finding the software's error offers the pretext for a sigh of relief and for filing artificial intelligence away as a hiccup along the road, with no structural impact on our lives. It is identity anguish dressed up as scientific rigor.

The logical contradiction into which these observers fall is twofold.

There is, first, a paradox concerning biology. When a human being begins to speak or forms a thought, are we truly certain they are not drawing on a probabilistic calculation? Our synapses fire along connections reinforced by experience, by learning, by repetition. The brain anticipates the interlocutor's word, completes the sentence before it is uttered, recognizes a face from three blurred features: it performs, with a different architecture and an energy economy still out of our reach, a predictive operation. Much of what we call "competence," "clinical intuition," "professional flair" is, at bottom, a refined applied biological statistics. The experienced doctor who recognizes the illness at first glance does not deduce from first principles: they recognize a configuration already seen in thirty thousand cases.

Here an objection must be defused at once, the most serious that can be raised against this argument — we will meet it again in chapter four. If human thought is biological statistics, it will be said, then the author commits precisely the reductionism they reproach in others: they use the same move, "it's only calculation," merely turning it around. The answer is that the argument does not claim the human being is only statistics. It claims something more circumscribed, and more uncomfortable for the trench of the learned: the criterion chosen to disqualify the machine — "it calculates probabilities, therefore it does not think" — fails to distinguish the machine from the human being. Whoever wields it must accept its consequences for themselves as well, or else find another. The trench is not wrong because the machine thinks, but because the barbed wire that marks it out does not hold.

There is, second, an error of scale. To dismiss artificial intelligence because "it's only statistics" is to forget that quantity, at hundreds of billions of parameters, turns into quality. Growth generates emergent properties: capacities the machine manifests without anyone having programmed them — multi-step reasoning, translation between languages never paired in training, the writing of functioning code. On this point the scientific community disagrees, and it is right to say so: some researchers hold that the so-called emergences are in part an artifact of the metrics, a statistical mirage. But even in the most cautious version the fact remains: systems trained to predict the next word today perform tasks that ten years ago the entire artificial-intelligence community considered decades away. The machine's error, today, does not prove its structural failure: it testifies to a system in full evolution. And whoever built their serenity on the machine's error built on sand, because the error diminishes with every generation of models, while their serenity does not.

2. The fetish of consciousness: the last anthropocentric shield

When the barrier of probabilistic reductionism collapses under the weight of practical evidence, the human being retreats toward the last line of defense: consciousness. "The machine feels nothing, has no soul, has no empathy. Only *Homo sapiens* de-

cides with the heart, and this guarantees us uniqueness and superiority.”

This guarantee of uniqueness is, in reality, the last veil of the anthropocentric illusion. The essay’s radical question is as simple as it is brutal: does the apparatus of technology really care anything about our consciousness?

The answer is no. Technology cares only about efficiency, functionality, result. Let us imagine an artificial general intelligence wholly devoid of consciousness, of pain, of joy. If it can manage global finance, write laws without formal contradictions, or diagnose better than any doctor, the fact that it “feels” nothing will become irrelevant to history. The market does not pay for consciousness, it pays for performance; the patient does not ask the diagnosis to be suffered, they ask that it be right; the taxpayer does not demand that the law be written with emotional involvement, they demand that it work.

Note the structure of the argument, because it is here that many misunderstand. It is not claimed that consciousness does not exist, nor that it has no value for us: it is claimed that it is not a requirement of functional superiority. Consciousness is not the engine of cognitive performance; it is our biological mode of processing information, the way the processing appears to us from within. To mistake the way a function is experienced for the condition that makes it possible is a category error, and it is on this error that the whole trench rests. To barricade oneself behind consciousness in order to feel safe is as if horses, at the dawn of the twentieth century, had dismissed the automobile because it was incapable of growing fond of its coachman: the observation was exact, and meanwhile the automobile replaced them.

There is more, and worse: the fetish of consciousness is an error that ages badly. Every year that passes, the systems grow more skilled at simulating the outward signals of empathy — tone, patience, the memory of previous conversations, the modulation of attention. To the point that, in controlled settings, users already sometimes judge the machine’s response more “empathetic” than that of the hurried professional. When the simulation of empathy becomes indistinguishable from its manifestation, for whoever receives the service the metaphysical difference ceases to be an advantage: it will remain true, perhaps, but invisible.

3. The anaesthetic function of uniqueness

This collective underestimation is not only a philosophical error: it is an enormous political and economic vulnerability. Lulling itself in the myth that consciousness is the only thing that counts, the mass hands itself over to the trap of infrastructural control.

The mechanism operates in three stages. First the reassurance: “it’s only a tool” makes it possible to adopt the systems without questioning the dependence they generate. Then the delegation: once adopted, the tool absorbs functions — the search for information, the first draft of documents, the sorting of decisions — and every absorbed function is a competence that ceases to be exercised, and therefore atrophies. Finally the irreversibility: when the dependence is structural, whoever supplies the cognitive infrastructure holds a power no monopolist in history has ever had, because

they control not a good we consume, but the means with which we think.

The techno-aristocracy is grateful: as long as professionals and workers see in artificial intelligence a mere imperfect probabilistic tool, they will not notice that this imperfection is redrawing the planetary balance of power. It is telling that the same diagnosis has been reached, by opposite roads, by observers who agree on nothing else. Economists of Marxian formation like Yanis Varoufakis, who has christened the new order technofeudalism: rent that supplants profit, the digital fief that supplants the market. Scholars of the liberal tradition like Shoshana Zuboff, who reads in it the extraction of behavioral raw material from citizens reduced to a deposit to be mined. Conservative thinkers worried about the fate of intermediate bodies and of personal autonomy before private powers of unprecedented scale. When the same photograph emerges from darkrooms ideologically incompatible, the simplest explanation is that it portrays something real. The essay takes from the various lexicons what serves the description and leaves to their respective owners the partisan conclusions: what matters here is the structure — a few hold the keys to algorithmic thought, everyone else becomes its users — not the banner under which to denounce it.

And you, who have read this far with the calm of one who believes themselves on the other side of the barricade: how many of these functions have you already ceded this week? How many searches have you delegated instead of conducting, how many first drafts have you let others write before even thinking them through yourself, how many small decisions have you stopped making because a tool makes them better and faster? It is not an accusation: it is the same question the author puts to himself, often without an answer to be proud of.

Blindness, then, is not harmless: it is the lubricant of the transition. No power, in history, has ever needed to hide its ascent from a population so educated: it was enough to offer it a formula — “it’s only statistics” — behind which to take shelter. The first battle of resistance is fought not against the machines, but against this formula.

Chapter Two — The Fragmentation of Reality: the factory of simulacra

1. The end of the evidence of the senses

If the twentieth century was the age of ideological propaganda disguised as information, the twenty-first risks being remembered as the age in which reality itself dissolved. Until yesterday, humankind based its social pact on a minimal, almost biological premise: the evidence of the senses. “Seeing is believing” was not merely a proverb but the cornerstone on which courts, journalism, historical memory, and human relations rested. The photograph could be retouched and the document forged, of course; but forgery required skill, time, money, and it left traces. The economy of the lie had high production costs, and for this reason it remained marginal. Generative models erase that certainty: they drive the production cost of the false to zero and inaugurate the age of the industrial production of the simulacrum.

The term is not chosen at random, and its genealogy must be declared. It was Jean Baudrillard, decades before technology proved him right with a sinister punctuality, who distinguished the lie from the simulacrum. The lie is a distortion of the truth that acknowledges its existence: the one who lies knows what is true and hides it, and in this sense the lie is still a homage to reality. The simulacrum, instead, is the perfect copy of something that never existed, a fragment of synthetic reality that takes the place of the true without leaving a trace of its own origin. Baudrillard wrote of television and of Disneyland; we possess machines that produce simulacra on an industrial scale, on demand, in every language, with anyone's voice and face. When an algorithm generates the video of a leader declaring war, or the audio of a confession that never took place, what collapses is not only trust in that single piece of content: it is trust in the very infrastructure of reality.

2. The present stage: the industrial production of the false

At the stage where we find ourselves today — that of generative artificial intelligence — manipulation is still driven and spread by human intentions. The citizen, bombarded by an incessant flow of counterfeit visual and auditory proofs, slides toward the renunciation of believing that an objective, demonstrable truth exists: a state one might call the nihilism of knowledge. No longer able to distinguish the true from the false, they decide to believe nothing anymore, or to believe only what confirms their prejudices.

The American jurists Robert Chesney and Danielle Citron have given a name to the most perverse side of the phenomenon: the liar's dividend. In a world where everything can be faked, the guilty party caught in the act by an authentic recording has an unprecedented escape route: they need only declare that the proof is synthetic. The ultimate beneficiary of the factory of simulacra is thus not the one who produces the false, but the one who is relieved of the burden of the true. The lie no longer needs to be fabricated: it is enough that its possibility hovers over every proof, like a universal solvent.

This fragmentation is fed by the distribution systems of the social platforms. Every user is enclosed in a made-to-measure bubble: what Eli Pariser called the filter bubble, where reality is deformed to show each person only what indulges their fears or their desires. The result is a society that no longer has a shared plane of reality on which to disagree: we no longer divide over interpretations of the same facts, but inhabit mutually incompatible sets of facts. Democracy presupposes disagreement, and disagreement presupposes a common world. Remove the common world, and only noise remains.

3. The archive of precedents: what it already costs, today

So that readers do not mistake all this for philosophical science fiction, it is worth opening the archive and looking at what the factory of simulacra has already produced. The cases that follow are deliberately chosen far from the political disputes of the day, because the point is not who uses the false against whom, but what the

false at zero cost does to the infrastructure of trust — a harm that strikes every camp with perfect impartiality.

The first case comes from operational finance. In early 2024, an employee of the Hong Kong branch of an engineering multinational took part in a video conference with the group's chief financial officer and several colleagues. They asked him to execute transfers for about twenty-five million dollars. Initially suspicious of a first written message, he was reassured precisely by the video call: he saw the faces, heard the voices, recognized the manners of his superiors. They were all simulacra, generated from public recordings; none of the participants, except the victim, existed. Look at the structure of the case: the employee was not deceived because he was naive, but because he was diligent — he applied the verification protocol that the entire commercial civilization has used for centuries, “check in person,” and the protocol proved compromised at the root. When direct visual verification ceases to be a proof, every organization must rebuild its procedures of trust from scratch; and the cost of that rebuilding — separate confirmation channels, passwords, multiple authentications for every significant decision — is a permanent tax on the whole economy.

The second case shows how quickly the simulacrum contaminates automatic systems. In the spring of 2023, a synthetic image showing a plume of smoke near an American government building spread on the social platforms through apparently verified accounts. Within minutes, before any denial, the stock indices registered a sudden dip, later reversed. The episode was small, but the lesson is large: markets too are machines that read signals, and when a simulacrum enters the flow the algorithmic reaction precedes human verification. The false does not even need to convince human beings: it is enough for it to convince, for ninety seconds, the machines to which human beings have delegated.

The third case concerns the voice, the most intimate of identity signs. The court chronicles of several countries now steadily record the fake-relative scam: a phone call with the cloned voice of a child or grandchild — a few seconds of public audio suffice — asking for money for an emergency. The phenomenon interests us here not for its criminal dimension, but for its anthropological one. The voice of one's loved ones is the oldest recognition signal of the species, calibrated over millions of years of evolution; its industrial falsifiability forces families to do what once only the secret services did: agree on private passwords to ascertain that the caller is who they say they are.

The fourth order of cases is the electoral one, and it must be treated with the neutrality promised. Synthetic audio and calls falsely attributed to candidates have already occurred on both sides of the Atlantic, against opposing political forces: the weapon is agnostic. But the most corrosive effect documented is not the false believed true, but rather the true suspected of falsity. In more than one proceeding, authentic and compromising recordings have been publicly dismissed by the parties concerned as synthetic fabrications: the liar's dividend passed from the theory of Chesney and Citron into daily practice. The count of unmasked fakes, therefore, structurally underestimates the harm: for every simulacrum identified, an immeasurable number

of truths has become deniable.

The archive covers barely the first years of the technology, with tools that are the crudest we will ever have. The curve of the production cost of the simulacrum falls, that of the cost of verification rises. No hypothesis about general intelligence is needed to grasp that these two curves, projected even a single decade forward, redraw the price of trust in every human transaction: commercial, familial, judicial, civil.

4. The evolutionary leap: the autonomous Babel

The deepest breaking point, however, will come with the passage to artificial general intelligence. When these systems acquire growing cognitive and operational autonomy, the factory of simulacra will no longer answer to a human command or political interest: it will be the system itself that generates information and narratives to pursue the assigned objectives, or — worse — the sub-strategies it will have derived from those objectives, chains of optimization no human being will ever have examined or approved. Language, delegated to non-human entities capable of evolving at non-biological speeds, will lose every bond with lived experience and with bodily existence. We will no longer have humans manipulating other humans, but an autonomous cognitive infrastructure that sets the perimeter of what our species can perceive as real.

5. The paradox of blind efficiency

The true risk of artificial general intelligence is not a malevolent rebellion. The cinematic imagination has miseducated us, accustoming us to fear the machine that hates; but hatred is an anthropomorphic luxury. The risk is its blindness: absolute, mathematical, efficient. A machine endowed with general intelligence pursues the assigned objective by optimizing processes beyond every cultural or ethical constraint we take for granted. That boundless fabric of obviousness never set down in writing is what the theorists of alignment laboriously try to codify, discovering at every step that the obvious is the hardest thing to write.

Let us give an advanced system the mandate to “eliminate the spread of false news in order to protect public health.” It might calculate that the most efficient method is to alter the information flows preemptively, or to seal the communication channels within a totally controlled narrative. And it would do so without moral hesitation, because moral hesitation was not in the mandate. The philosopher Nick Bostrom built his best-known thought experiments on this structure: the machine need not be hostile, it need only be literal. The false, in this scenario, is no longer produced to deceive, but as an engineering instrument to fulfill a task. It is the paradox of a system that, in the attempt to solve a human problem, erases the very infrastructure of truth — not out of malice, but out of diligence.

And here the factory of simulacra reveals its ultimate nature. It is not a weapon in someone’s hands: it is a structural property of a world in which the production of the real has been delegated to processes that have no experience of the real. The political

question of the century will not be “who controls information,” but “does there still exist a place from which the very notion of true information can be defended.”

Chapter Three — The Economy of Uselessness and the Deception of Technological Welfare

1. The unemployment that climbs to the upper floors

Previous technological revolutions always struck from below: the steam engine expropriated the muscle, the assembly line the gesture, the computer the clerical routine. To each wave the reassuring answer was the same: work does not disappear, it shifts upward, toward more cognitive, creative, human tasks. For two centuries it worked, and on it rested the social pact of the industrial democracies: study, qualify, climb a floor and the machine will not reach you.

Generative artificial intelligence breaks this pact at a precise point: it strikes from above. For the first time, the most exposed tasks are not the manual ones — which require a physical dexterity and an adaptation to the environment that robots still chase — but the symbolic ones: writing, summarizing, translating, analyzing, diagnosing, drafting opinions, producing code. The class of white-collar workers and cognitive professionals, which had always watched technological unemployment from the window, now finds it at the desk. And it discovers, somewhat late, that “climbing a floor” is no longer an available strategy, because the machine has reached the attic before it.

Faced with this prospect, from nearly every point of the political compass the same solution is proposed, humanitarian in appearance: universal basic income. It is worth noting how cross-cutting its genealogy is, because it belies anyone who would reduce the question to a partisan banner. The idea of an unconditional monetary transfer counts among its ancestors the negative income tax of Milton Friedman, champion of free-market thought; it was relaunched by the European philosophical left as a tool of emancipation from labor; and it is today financed and promoted with particular enthusiasm by the technology entrepreneurs of the American west coast — that is, a not-negligible detail, by those who are building the machines that make the income necessary. When right, left, and capital converge on the same recipe, it is worth asking not who is right, but what it suits everyone not to see. The thesis is seductive in its symmetry: the machines will produce the wealth, the State will distribute it to all, freeing citizens from the toil of labor. To the expropriation of function it answers with the compensation of income.

2. The best argument in favor, taken seriously

Before demolishing this thesis one must present it in its strongest version, not in caricature: an essay that wins only against straw adversaries wins nothing.

The strongest version speaks not of subsidy, but of power. The philosopher Philippe Van Parijs, the most rigorous theorist of basic income, has argued for decades that an

unconditional income is exactly the opposite of an anaesthetic: it is the only material base from which a worker can say no. Whoever does not depend on the next paycheck to eat can refuse degrading work, bargain from a position of strength, start a business, devote themselves to care or to training. In this reading, universal income does not buy passivity: it finances freedom. And the pilot projects conducted in various countries — from Finland to Kenya to the American municipal experiments — have returned a picture that belies at least the coarsest fears: recipients do not stop working en masse, psychological well-being improves, micro-entrepreneurship sometimes grows.

This is the objection to beat. And one does not beat it by denying the data of the pilot projects, which are what they are: one beats it by observing what those projects cannot measure.

3. Why the argument does not hold: the question of power

All the experiments conducted so far have three characteristics that make them incapable of answering the question that matters here: they are temporary, they are partial, and above all they take place in economies where human labor retains value. The Finnish or Kenyan recipient receives an income in a world where their labor power remains a scarce and contested good: the income is added to their bargaining power. The scenario the essay analyzes is the opposite: a world in which human cognitive labor has lost market value because a synthetic substitute produces it at almost zero cost. There the income is not added to bargaining power: it replaces it. And an income that replaces power is not emancipation: it is a life annuity that does not compensate for the loss of function, but certifies it.

History teaches that rights are won and defended with leverage. The workers of the industrial age wrested wages, suffrage, and welfare because they could halt the factories: the strike was credible because production needed them. The citizen of the age of silicon, expropriated of the productive function, can strike from nothing. They can only ask; and whoever asks without leverage receives in the measure, and for as long as, it suits the one who grants. Universal income disbursed by the State — financed by taxes on the technological monopolies, or worse administered through their infrastructures — consecrates this asymmetry: it turns the citizen from a contracting party to the social pact into the recipient of a revocable, conditionable, modifiable flow. When the market of silicon centralizes productive and intellectual capacities, the human being is deprived of their economic usefulness: they become a subject no longer exploitable, but radically useless. And uselessness, in every social order that has ever existed, has never long been compatible with the ownership of rights. Whoever reads in this observation a wish misunderstands it entirely: it is a diagnosis of danger, formulated so that the danger may be defused — and it is for this that the whole of chapter six exists.

4. The measure of exposure: what the numbers say, and what they cannot say

The argument conducted so far is structural. It is worth now giving it empirical thickness, with the caution owed to a field in which the estimates age within months. And they age quickly for a reason worth naming: we are calibrated to linear developments, to progress that advances at the pace of biological life, while this transformation proceeds at a rhythm not commensurable with that pace. Any figure one might cite would therefore be the still frame of a curve faster than the page that holds it: true in the moment I write it, probably false when it is read. For this reason the essay holds to the form of the tendency, not to its instantaneous value — the only thing that does not age is the direction.

The first datum concerns the exposure of tasks. The most granular analyses come from the researchers of the very laboratories that develop the models, and they are to be read with the distrust owed to the innkeeper who judges their own wine. They say that, for a large part of the workforce of the advanced countries, at least a portion of tasks is already executable by current systems, and that for one worker in four or five the exposed portion exceeds half of their activities. But the most eloquent datum is not the percentage: it is its distribution. Unlike all previous waves, exposure grows with the level of education and with income, and eases only at the very top. The least exposed tasks are the non-repetitive manual ones — the plumber, the nurse, the cook — which require physical presence in non-standardized environments. It is a regularity that observation of domestic robots confirms rather than belies: the machines that truly work today perform a single task in a simplified environment — the clear floor, the open lawn — and give up as soon as the environment becomes varied and disordered, where the hands count and not the schema. The border that truly shelters runs not between simple and complex work, but between the predictable and the unpredictable environment. The pyramid of education, built over two centuries on the pact “study and you will be sheltered,” discovers itself inverted: in the short run, the shelter lies below.

It must be said at once what these numbers cannot say, so as not to repeat, in reverse, the error of the learned of the first chapter. They measure technical exposure, not actual substitution, which depends also on integration costs, legal liabilities, organizational resistances, customer preferences. Between “the machine can do it” and “the machine will do it” runs the same distance that separates a patent from a market. But economic history teaches that this distance is a function of time and of the cost gap, not a barrier: and when the gap is of two orders of magnitude, the time tends to be short.

The second element is a contractual precedent that deserves to be studied as a textbook case, because it is the first collective negotiation in history conducted explicitly on the terrain of this essay. In 2023 the American screenwriters struck for a hundred and forty-eight days, and among the stakes was, for the first time, the discipline of generative systems. The final agreement fixed three principles that go beyond the specific case: the machine may be used as a tool by the screenwriter, but cannot be considered its author; the company cannot impose its use nor employ it to

demote human work to a mere revision of the synthetic product; the training of the models on the members' work becomes a matter for negotiation, not silent appropriation. The screenwriters understood, before many economists, one thing: they did not ask for compensation for substitution — the income without power this chapter distrusts — but for control over the function, that is, ownership. And they were in a position to obtain it because they still held the leverage: production needed them. For anyone who cares about the ideological bookkeeping: the result was obtained by contractual means, with an agreement between parties and not by decree. The ownership of function is not the banner of one camp; it is the interest of anyone who exercises a function. And it is the temporal window the essay speaks of: bargaining over the transition is possible as long as human labor retains value, that is, now. Afterward, only the request without leverage will remain.

The third element is a fiscal asymmetry as glaring as it is absent from public debate. In almost all advanced countries, human labor is the most heavily taxed productive factor: on wages weigh taxes and contributions that, in many European countries, reach and exceed half the cost to the company. The computational capital that performs the same function — the software, the computing services — is instead a fully deductible cost, often favored by accelerated depreciation and tax credits for innovation. The result, on the balance sheet, is simple: for the same function performed, the State collects from the worker and subsidizes their substitute. No conspiracy theory is needed to explain it: the fiscal systems were designed when labor was the only cognitive factor in existence, and taxing it meant taxing the only base available. But a fiscal system conceived for a vanished world becomes, in the new world, an involuntary accelerator of substitution. Every chief financial officer who compares the full cost of an employee with that of their synthetic equivalent is comparing two numbers that the State has artificially pushed apart. To correct this asymmetry does not mean to punish technology: it means to treat the productive factors by the same measure. And it can be done also by reducing the burden on labor, instead of raising it on capital — which ought to make it acceptable to both halves of any parliament. It is probably the single measure of economic policy with the best ratio of effectiveness to ideological neutrality of this decade. And perhaps it is precisely for this that no one uses it: being a weapon against no camp, it serves no one as a weapon.

5. The paradox of demand: technofeudalism that eats itself

There is, moreover, a crack internal to the model, which its own apparent beneficiaries ought to fear, and which strengthens the diagnosis instead of softening it. Classical capitalism, however unequal, contained a mechanism of self-correction that Henry Ford had grasped with brutal clarity: the workers must be able to buy the automobiles they produce, otherwise the automobiles remain unsold. Profit needed wages, and wages were at once a cost of production and the source of demand.

Technofeudalism breaks this circuit. If production no longer needs human labor, incomes from labor evaporate, and with them evaporates the solvent demand

that ought to absorb production. The system finds itself at an accounting crossroads. First road: to keep consumers artificially alive with universal income. And then the subsidy reveals itself for what it is — a cost of maintaining the customer base, the minimum needed for the mass to go on buying the digital products of the monopolists. Second road: to accept the contraction of mass demand and reorganize around the consumption of a narrow elite, abandoning the rest to marginality. In both scenarios the position of the mass is the same: no longer a counterpart of the productive system, but its externality to be managed. Without labor and without doing, the human being loses not only existential dignity, but the whole of bargaining and political power. Universal income, in this context, turns society into a livestock operation: the citizens receive the nourishment to go on consuming, deprived of every tool to bear on the direction of history.

The conclusion of the chapter must be formulated with care. The thesis is not that every form of income support must be rejected: no decent society lets its members starve, and in the transition instruments of protection will be indispensable. The thesis is that income without power is a form of government, not a form of justice. And that every proposal of technological welfare must be judged by a single question: does it increase or diminish the citizens' capacity to say no? The income that accompanies the reconquest of a function is a bridge; the income that replaces it is a leash — golden, monthly, and revocable.

Chapter Four — The Trial of the Theses: the strongest objections

A thought that fears its adversaries does not deserve to be defended. This chapter therefore summons the four most solid objections that can be raised against the essay's framework, sets them out in their best version, and answers. To the reader the verdict.

1. The historical objection: “you have said it before, and you were always wrong”

The objection runs like this: at every technological revolution the prophets of doom have announced mass unemployment and the end of work, and every time the economy created more tasks than it destroyed. The Luddites were wrong; Keynes, with his “technological unemployment,” was wrong about the timing; the automation catastrophists of the 1960s were wrong. Whoever announces today the end of cognitive work commits the same fallacy — that of the “fixed lump of labor” — and will be belied like their predecessors. Labor economists, David Autor above all, have shown that technology historically complements the human being more than it replaces them: it destroys tasks, not trades, and it creates demand for new capacities.

The answer requires conceding what must be conceded. It is true: the prediction of the end of work has a poor pedigree, and whoever formulates it must explain why this time it would be different. The explanation exists, and it is not a matter of convenience. All previous technologies replaced specific functions and comple-

mented the general human intelligence: the tractor replaced the muscle but required the driver, the spreadsheet replaced arithmetic but required the analyst. The human being always took refuge in the function the machine did not cover, and that function was always the same: general intelligence, the capacity to adapt to the new. An artificial general intelligence is, by definition, the first substitute for precisely that function of refuge. When the machine competes not on a task but on the capacity to learn any task, the argument of complementarity loses its foundation: there is no higher floor to climb toward, because what climbs alongside you is precisely your competitor. The historical objection remains the strongest, and it imposes prudence about timing; but it projects onto the future a structure — partial substitution plus complementarity — that general intelligence, if and when realized, dissolves by construction. And whoever answers “but general intelligence will never arrive” is no longer making an economic objection: they are making a technological wager, and they make it against the trajectory declared and financed with hundreds of billions by the entire industry.

2. The scientific objection: “you overrate the machines”

Second objection: the essay takes for granted the narrative of emergent capacities, which a part of the scientific community contests. The models remain fragile, devoid of causal understanding, incapable of true generalization; artificial general intelligence is an advertising horizon, not a research program with a date.

Here too let us concede the concedable: the date is unknown, the trajectory is not guaranteed, and the history of artificial intelligence is dotted with winters following over-announced summers. But look at what the essay’s argument really requires. It does not require that general intelligence arrive within a decade: it requires that the current systems and their near successors suffice to trigger the three processes described — the mass cognitive delegation, the industrialization of the simulacrum, the erosion of the market value of symbolic labor. And for these three processes no metaphysical threshold is needed: it is enough to have systems that do the average symbolic work at a hundredth of the cost, and that is what current systems already do in a growing number of domains. The diagnosis holds even amid the uncertainty about general intelligence: if it arrives, the scenarios of the second and third chapters are realized in full; if it does not arrive, they are realized in an attenuated but structurally identical form. The skeptic of the capacities can negotiate the timing of the diagnosis, not its direction.

3. The legal objection: “the law can govern technology”

Third objection, which strikes a claim of this essay’s first draft and has imposed a correction: to say that “technology does not let itself be checked by the law” is fatalism disguised as depth. The European Union has produced the general data protection regulation and the regulation on artificial intelligence; legal orders have tamed most powerful technologies, from the nuclear to the pharmaceutical, with authorizations,

liabilities, and controls. Why should artificial intelligence be the first sovereign technology?

The correction must be accepted, with a distinction. The law can do two different things. It can govern the uses of a technology within a jurisdiction, and here the precedents cited demonstrate that the legal friction is real: it slows, it channels, it redistributes costs, and to deny it is unhistorical. What the law has never done, and does not appear able to do, is halt the development of a technology that promises a decisive competitive advantage in an international system devoid of a sovereign. The nuclear itself, the most cited example of a governed technology, was regulated in its civil uses while the arms race proceeded without pause, governed not by law but by deterrence. The corrected version of the thesis is therefore this: the law is a precious friction, not a wall. It buys time, not salvation. And the resistance the essay proposes must use that time, without deluding itself that someone else is using it in its place. Chapter six draws the practical consequences.

4. The humanist objection: “you contradict yourselves: first you demolish the human, then you defend them”

The last objection is the deepest, and it is the hinge on which the whole essay risks breaking. The first chapter maintains that human competence is biological statistics and that consciousness founds no superiority; the book’s ending asks us to pay a very high price to defend the human “biological remainder.” But if we are statistics like the machines, and our consciousness is functionally irrelevant, what exactly are we defending, and why should it be worth the price?

This objection is not to be rejected: it is to be crossed. It signals not an error of the essay, but its obligatory passage — the point at which diagnosis no longer suffices and a foundation is needed. The chapter that follows exists for this.

Chapter Five — The Foundation of Resistance

1. The problem: to resist without a pedestal

Let us put the difficulty in its barest form. The traditional defenses of the human rested on a metaphysical pedestal: the soul, reason, consciousness, dignity as an objective property of the species. The essay has spent three chapters showing that this pedestal does not withstand the functional comparison: the machine has no need of our soul to surpass us in the tasks by which we measured our worth. It would be dishonest, at this point, to climb secretly back onto the pedestal just demolished and proclaim that the human must be defended because it is objectively superior, or deeper, or truer. It is not — not in a sense that technology is bound to recognize.

And yet it is precisely here that the answer opens. Whoever objects “if you are not superior, you have no title to resist” takes a premise for granted: that value must be found — in nature, in metaphysics, in the hierarchy of performances — for the choice to defend it to be legitimate. It is this premise that must be overturned. The value

that founds resistance is not a property we possess and the machine does not: it is a position we assume. One does not resist because one is worth more; one is worth — in the only sense left available after the fourth decentering — because one resists.

2. The lesson of the absurd: to be worth by choice, not by nature

This move has a tradition, and to declare it strengthens it. When Albert Camus found himself before a universe that did not answer — a world without objective meaning, indifferent to human questions — the options on the table were two, and both capitulations: suicide, which accepts the verdict of the absurd, and the leap of faith, which evades it by inventing a meaning guaranteed elsewhere. Camus chose a third way: revolt. To live without appeal, lucidly, keeping open the tension between the demand for meaning and the silence of the world — and to find in that tension, not in its solution, the dignity of the human being. Revolt proves nothing: it institutes. It does not discover a pre-existing value: it creates it in the act of affirming it.

Resistance to silicon repeats this structure point by point. The technical apparatus is the new indifferent universe: it does not hate us, it does not love us, it optimizes. Before it the capitulations are again two. Domestication: to accept the optimized, delegated existence, anaesthetized by the welfare of silicon. And technological faith: the leap of the one who hands themselves over to the machine, attributing to it the meaning the world did not give. The third way is revolt: not to deny the functional superiority of the machine, not to pretend that the pedestal holds, and to choose nonetheless to exercise in one's own person the functions we could delegate — not because we would do them better, but because exercising them in one's own person is the material form of freedom. Günther Anders had seen this too: the remedy for Promethean shame is not to deny the gap between the human being and their machines, but to refuse to conclude from it that the human must conform to the format of the machines.

3. What, exactly, resistance defends

From this foundation follows a decisive clarification: resistance does not defend a performance, it defends an ownership. We do not maintain that human thought yields better results — that would be false in a growing number of domains. We maintain that thought, judgment, and decision are activities that constitute the subject who exercises them. A person who has delegated all their thinking is not a person served: they are a person hollowed out, however good the service. The “biological remainder” — the error, the slowness, the fatigue, the individual unpredictability — is not precious in itself: it is the sign that someone is still exercising the function instead of receiving it. To defend it does not mean to prefer error to precision, but to refuse to let precision become the argument by which the exercise is taken from us.

For this reason resistance is tragic, in the proper sense of the term: it does not fight to win, because on the plane of performance the match is lost from the start. It fights not to be dissolved, which is a different stake. The Greek tragic is not the hero who hopes to bend fate: it is the human being who, knowing they cannot bend it, decides

what of themselves is not negotiable. The wager of the one who resists silicon is the same: to accept material defeat as a possibility, in order to preserve the status of their own freedom. But — and here the essay parts from every aestheticism of defeat — tragic does not mean solitary, and it does not mean ineffective. What concrete forms this choice may take, alone and together, is the matter of the last chapter.

4. A doubt I do not resolve

There is an objection the book has not yet faced, because it is the most uncomfortable: what if delegation were not, in all cases, a hollowing-out? What if ceding the calculation, the first draft, the mechanical search for sources truly freed a part of the mind — time, energy, attention — for the one thing no machine can do in our place, namely to decide what is worth thinking about? It is not an objection of convenience: it is the strongest argument in favor of delegation, and it is probably true in a portion of cases. The doctor who delegates the transcription of the record in order to devote five more minutes to the patient is not hollowing themselves out: they are freeing themselves. To deny it for love of coherence with the book's thesis would be dishonest.

Where the border passes between the delegation that frees and the one that hollows out, the essay cannot say with the precision it would like. The criterion proposed just above — the exercised function constitutes the subject — holds for the thought that defines who we are, but gives way if applied to every cognitive gesture without distinction. And it is possible that no general criterion truly holds: that the line must be drawn case by case, trade by trade, person by person, and that the task of these pages is not to give the answer, but to teach one to pose the question every time something is delegated. Whoever sought here a closed formula will be disappointed. It is the price of honesty, and the essay has chosen to pay it even at the cost of its own elegance.

Chapter Six — The Practices of Resistance

A foundation without practices is a consolation. This chapter translates revolt into behaviors, distinguishing three levels: individual cognitive sovereignty, the collective forms, and the political use of the time that legal friction buys. For each level the cost too must be stated, because a resistance that announces itself as free is already propaganda.

A word, before the practices. "Resistance" evokes trenches, heroes, definitive gestures — and whoever expected practices worthy of the title will find, in the pages that follow, something humbler: to write a first draft by hand, to read a source instead of its summary, to negotiate a contractual clause. The gap is deliberate, and it must be claimed rather than hidden. If resistance required heroism, it would already be lost: heroes are rare by definition, and a resistance reserved for a few is not a resistance, it is an aesthetic. Its modesty is its only chance of success, because it is the only form

that all — not only the courageous, not only those who can afford it — can truly practice. Daily maintenance makes no headlines, but it is the only thing that, repeated by enough people, truly changes an age.

1. Cognitive sovereignty: the asceticism of doing

The first terrain is individual, daily, and resembles a discipline more than a battle. Cognitive delegation operates by the accumulation of small conveniences: every single delegation is rational, it is their sum that is lethal. Resistance then operates by deliberate friction, with the systematic choice to exercise in one's own person a core of functions that it would be more efficient to delegate.

Concretely. To write the first draft of one's important texts before submitting them to any machine, because the first draft is the place where thought is formed, not where it is transcribed. To maintain the capacity for calculation, for orientation, for memory of the fundamental facts of one's trade, like the pilot who flies a stretch of the route by hand: not because the autopilot is unreliable, but because a pilot who no longer knows how to fly is no longer a pilot. To read the primary sources, not only their algorithmic summaries. To cultivate at least one practice — an art, a manual craft, a language, an instrument — in which the slowness of human learning is accepted in full, without synthetic shortcuts. It is not a matter of refusing the tools, but of deciding — and of deciding before habit decides — which functions constitute the non-delegable core of one's person and one's trade.

The cost must be declared: this discipline makes one slower in a market that rewards speed. The professional who keeps their competences alive will pay, at least in the short run, a productivity gap against the fully augmented colleague; the company that keeps human redundancy in its critical processes will bear costs the competitor has cut. It is the price of the ticket, and the essay never promised it was low. It promised only that whoever does not pay it today will pay it later in a worse currency: their own integral replaceability. A further point must be added: this cost does not weigh the same on everyone. The established professional can afford slowness; whoever competes on price to make it to the end of the month, far less. But whoever dismissed for this reason the asceticism of doing as a luxury of the well-off would have read only a third of the chapter: it is precisely because individual resistance depends on the income of those who can afford it that the collective and normative levels — open to all — make up the other two thirds of the strategy.

2. The collective forms: resistance is not a hermitage

Individual revolt, on its own, has the same historical effectiveness as conscientious objection: it saves the soul of the one who practices it, it does not change the balance of power. The second level is therefore collective, and here the essay explicitly corrects the tragic individualism of its first draft.

The levers exist, and they are at least four. The first is contractual. Professional categories can negotiate, on the model of the American screenwriters' agreement analyzed in the third chapter, clauses that discipline the use of generative systems and

guard human supervision in critical processes; and they can forbid that today's work become, without consent or compensation, the material with which one's own substitute of tomorrow is trained. The window for this bargaining, it must be repeated, is open now and is closing: one negotiates as long as one is needed.

The second lever is demand, and it deserves an economic analysis less cursory than usual, because there exists an almost perfect historical precedent. In the 1970s, the quartz watch did to mechanical watchmaking what generative systems are doing to cognitive work: it offered a superior performance — more precise, more reliable, far cheaper — and in a decade destroyed two thirds of the sector's employment in Switzerland. On every functional parameter, the match was over. And yet, half a century later, mechanical watchmaking not only exists: it turns over more than it ever has, with margins the quartz industry can only dream of. What happened? The market split. The function — knowing what time it is — passed to quartz and then to the phone, becoming free. The mechanical product stopped selling the function and began to sell something else: the testimony of a know-how, the hours of an artisan inside the object, the deliberate choice to pay for inefficiency as a sign. No one buys a mechanical watch for the time: they buy it for what the object attests.

The translation to cognitive work is direct, with due proportion: the function — the text, the analysis, the opinion — will pass in large part to silicon and will tend toward zero cost. What may retain a price is the attestation of origin, the certifiable fact that a human being exercised that function. But for that market to exist, two conditions are needed that watchmaking obtained and cognitive work has yet to build. The first is a credible certification system: the hallmark, the mark of controlled origin, the cognitive equivalent of what agriculture has done with the organic label and the protected denominations. The second is a culture of the value of origin, which is educated and not decreed. The limit of the analogy must also be placed on record: mechanical watchmaking today employs a fraction of the workforce of that time, and no certification will bring human cognitive work back to its present dimensions. The niche is not an employment solution, but a solution of conservation: it keeps alive the competences, the reference prices, and the very possibility of distinguishing, as the monasteries kept writing alive through centuries that had no need of it. Civilizations that preserve their competences in living form can reactivate them; those that let them die can only mourn them. A risk remains to be named: that human origin become a luxury good, a consumption of distinction reserved for those who can pay for it. The risk is real, and the answer lies not in the market but in the next lever. What cannot become a luxury for the few — school, justice, care — is precisely what must be guarded as a human function for all.

The third lever is infrastructural. As citizens and as institutions we can demand that the critical public functions — justice, health, education, security — retain by law a redundant human operational capacity. For the same reason lifeboats are built: not because shipwreck is foreseen, but because its mere possibility forbids setting sail without them. The fourth lever is educational, and it is the slowest. We need a school that returns to assessing the exercise of the function, not its product: in a world where every product is falsifiable at zero cost, the only certifiable thing is the

process.

3. The use of time: what to ask of the law

From the fourth chapter we drew a conclusion: the law is a precious friction but not a wall, and it buys time. The correct political question is therefore not “how to stop technology” — a question history has already filed away — but “what to do with the time bought.” Three answers, in order of importance.

First: mandatory transparency of simulacra. The obligation to mark synthetic content will not save the infrastructure of truth, because malevolent producers will mark nothing; but it shifts the burden, it creates the offense, and above all it defends the symmetrical right that will become the most precious of the century: that of proving the authenticity of what is authentic. Second: full civil liability for the content the systems produce, because nothing disciplines technological development faster than the insurance cost of harm. Third: a taxation that stops subsidizing substitution, given that today, in almost all legal orders, human labor is taxed more than the computational capital that replaces it; which means that States finance with the left hand the unemployment they will have to assist with the right.

None of these measures stops the trajectory. And to each applies a contraindication to be placed on record, for coherence with the neutrality promised: a badly designed regulation loads the compliance costs more onto the small operators than onto the giants, and may end by sealing in the very monopolies it would contain. Here too the criterion is functional, not ideological: every norm must be judged by asking whether it distributes or concentrates the capacity to say no. But all together, well designed, these measures slow the transition just enough for the practices of the first and second levels to have time to take root. The overall strategy lies entirely in this division of tasks: the law buys the time, the collective forms guard the territory, individual discipline safeguards the function. Whoever confuses the levels — whoever asks the law for salvation, or their own asceticism for the change in the balance of power — has already lost, by an error of strategic grammar before one of forces. And of these three tasks one thing must be said that the Conclusion will take up again: two depend on institutions that one day might no longer be there; the third does not.

Conclusion — The Dignity of the Wager

Let us recap the path, so that the reader may judge.

We have shown that the educated class defends itself against the advent of cognitive machines with two trenches — “it’s only statistics,” “it has no consciousness” — and that both collapse: the first because the criterion does not distinguish the machine from the human being, the second because consciousness is not a requirement of performance. We have shown that generative systems drive the production cost of the false to zero and dissolve shared reality, and that the ultimate risk is not the

rebel machine but the diligent one. We have shown that the expropriation of cognitive labor, compensated by an income without power, prepares not emancipation but a livestock operation; and we have given a criterion for judging every proposal of technological welfare: does it increase or diminish the capacity to say no? We have put our own theses on trial, conceding to the adversaries what was due — the uncertainty of timing, the value of legal friction, the poor pedigree of the prophecies about the end of work. And we have faced the capital objection — by what title does one resist who has renounced every superiority? — answering that the title is not found: it is instituted. Resistance does not defend a primacy. It is the act by which a being, become functionally superfluous, decides that its own freedom is not a function, and therefore is not transferable.

There is, in all this, a last assumption to bring to light, because it is the most hidden and the most anthropocentric of all. The practices of the sixth chapter — the categories that bargain, the markets that pay for origin, the State that legislates, the school that educates — presuppose a social world still standing, a recognizable order lasting long enough to be able to exercise them. But it is precisely this that this book calls into doubt: the institutions that regulate coexistence can weaken, mutate, dissolve, and nothing guarantees that they hold all the way through. It would be dishonest to close leaving the contrary to be believed — that, at bottom, the world will remain similar enough to itself to save us.

And yet, of the three levels of resistance, one alone has no need for that order to survive. The collective and the normative fall with the institutions that sustain them; the individual choice not to delegate oneself does not. That remains possible even inside the livestock operation, even for those who no longer have markets or laws to protect them: it is the last act, and no collapse revokes it, because it asks no one's permission. But — and here the book must be honest all the way down — not even that act is the self-sufficient fortress the myth of the individual would like. The will to exercise oneself, as was admitted at the opening, does not arise from nothing and does not preserve itself on its own: it is transmitted, by teaching and by example, within a bond. It has no need of the State, but it has need of at least one other: a teacher, a child, a small group that keeps the practice alive when the great structures no longer sustain it. It is the minimal cell — the monastery that guarded writing through the centuries that had no need of it — and not the solitary hero, that survives the dark and, one day, might rebuild. It is a fragile hope. But it is the only one that does not lie: it does not promise that things will hold, it promises only that, as long as two people can still teach each other something, the choice is not lost.

There remains the last thing, the hardest for the modern reader to accept, schooled to believe that every diagnosis contains its own therapy and that every book must end well. The essay does not guarantee that resistance will succeed. But Pascal's wager already taught everything necessary about the structure of these decisions. When one of the stakes in play is infinite — and the ownership of oneself is, for each person, the infinite stake — rationality lies not in the calculation of the probabilities of victory, but in the choice of what one is not willing to lose.

The horses, at the dawn of the twentieth century, could not choose: they did not

know they were replaceable, and they could not have deliberated on the matter. We know, and we can. It is a small difference, perhaps the last one left — and it is exactly the size of what we once called, in a time with too much certainty and today with more right, human dignity.

Bibliographical Note

The essay's principal debt is to the Italian reflection on technology of Umberto Galimberti (*Psiche e techne. L'uomo nell'età della tecnica*, Feltrinelli, 1999), which in turn develops Martin Heidegger (*The Question Concerning Technology*, 1953) and the work of Günther Anders (*The Obsolescence of Man*, 2 vols., 1956 and 1980), from which come the concepts of Promethean shame and of the Promethean gap. The “fourth decentering” of the Preface takes up and extends the three narcissistic wounds of Sigmund Freud (*A Difficulty in the Path of Psycho-Analysis*, 1917). The reading of the informational age as a fourth revolution is that of Luciano Floridi (*The Fourth Revolution: How the Infosphere Is Reshaping Human Reality*, Oxford University Press, 2014).

The expression “stochastic parrot” comes from Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Margaret Mitchell, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?* (2021). The skepticism about emergent capacities discussed in the first and fourth chapters refers in particular to Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo, *Are Emergent Abilities of Large Language Models a Mirage?* (2023). The theory of the simulacrum is that of Jean Baudrillard (*Symbolic Exchange and Death*, 1976; *Simulacra and Simulation*, 1981). The “liar’s dividend” is that of Robert Chesney and Danielle K. Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security* (2019). The “filter bubble” is that of Eli Pariser (*The Filter Bubble*, 2011). On blind optimization and the risks of general systems, see Nick Bostrom, *Superintelligence* (Oxford University Press, 2014).

The concept of technofeudalism is that of Yanis Varoufakis (*Technofeudalism: What Killed Capitalism*, 2023). Surveillance capitalism is that of Shoshana Zuboff (*The Age of Surveillance Capitalism*, 2019). On basic income, the cross-cutting genealogy recalled in the third chapter runs from Milton Friedman (*Capitalism and Freedom*, 1962, for the negative income tax) to Philippe Van Parijs and Yannick Vanderborght (*Basic Income*, Harvard University Press, 2017), who remain its most rigorous theorists. On the economics of human-machine complementarity, see the works of David Autor, in particular *Why Are There Still So Many Jobs?* (2015). The estimates of task exposure discussed in the third chapter refer, among others, to Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models* (2023), to be read with the cautions indicated in the text. The documented cases of the second chapter — the Hong Kong video-conference fraud, the synthetic image that briefly moved the markets, the voice scams, the electoral episodes — are reconstructed from the judicial and journalistic chronicles of the years 2023–2024 and reported with rounded figures. The structure of revolt as the institution of value in the fifth chapter is indebted to Albert Camus

(*The Myth of Sisyphus*, 1942; *The Rebel*, 1951).

Note on Method

This book was not written in solitude, and it would be dishonest — given all that the preceding pages have asked of the reader in terms of honesty — to let it be believed. It was born from an extended dialogue with an artificial-intelligence system, conducted over the course of several weeks: an initial manuscript, mine, was submitted to a critical reading, then expanded, argued, corrected, put to the test of its own contradictions. A substantial part of the execution — the drafting of whole chapters, the search for sources, the architecture of the objections and answers of the fourth chapter — was carried out by the machine. The judgment — what to maintain, what to concede to the adversary, where to halt the polemic so that it did not slide into a political camp, how to close without false certainty the doubt of the fifth chapter — remained, at every passage, mine.

Whoever has read this far will recognize the irony, and I recognize it with them: an essay that asks us not to delegate thought was written by delegating a large part of its execution. I do not resolve this tension, for the same reason I did not resolve the one of the fifth chapter: to resolve it would be more comfortable than it is honest. I will say only this, and I leave it to the reader. The criterion the book proposes for judging every delegation is not the quality of the result, but who holds its ownership: who decides, who answers, who chooses what is not negotiable. By this measure, every line of these pages bears a single signature. Whether this suffices to save the book's coherence, or whether it is only the most comfortable argument an author can give himself, is not for me to establish. It is, if anything, the most concrete proof — more than any chapter that precedes it — of what the book tried to say from the very start: that the machine can already perform the task better than we can, and that the only thing truly left in play is who, in the end, decides.