

# *Icarus:*

## Disembodied Knowledge, Bureaucratic Thinking, and the Hopeful Return to Reality

Dr James Wilk  
*University of Oxford*

Neither will he calmly transfer to the machine made in his own image the responsibility for his choice of good and evil, without continuing to accept a full responsibility for that choice.

—Norbert Wiener, *The Human Use of Human Beings* (1950)

We need a revolution within ourselves. ...  
We need to radically alter the way we conduct our lives.  
We need to change the way we relate to each other.  
We need to rearrange our goals and desires. ...  
I mean it. It is no joke. It is no game. ...  
Commit. Agitate. Educate. Speak. Act. Learn. Disobey.  
Rage against the machine.

—Kent McClard, in *No Answers, No. 9* (1989)

...a risky and hopeful return to a position before  
discouragement and defeat had gained the upper hand...

—Peter Lomas, *True and False Experience* (1973)

All action is premature until it's too late.  
—David Bernstein, 1983

### *Part I: A Statement of the Problem*<sup>i</sup>

#### **Two myths**

For nearly four centuries the Western mind has been labouring increasingly under a crippling limitation, contrary to the spirit and tradition of the previous 2000 years of Western thought, and so inimical to all that's best in European culture—an affliction sapping our creativity and marginalizing common sense. Often characterized as a collective madness, it has become, sadly, our biggest and most dubious export to the emerging markets.

I'm referring to the intellectual syndrome known as Rationalism, which I'll briefly describe in a moment, and which the philosopher Professor Michael Oakeshott rightly called “the most remarkable intellectual fashion of post-Renaissance Europe.” Rationalism has had far abler, more distinguished, and more articulate critics than me, and in increasing numbers in recent years. The most important and best known within the anglophone tradition have been the

philosophers R.G. Collingwood, Ludwig Wittgenstein, Stephen Toulmin, and Oakeshott himself, but distinguished European critics on the Continent are legion.

Yet this persistent syndrome, Rationalism, goes from strength to strength, propelled by its own hubristic “presumption of inevitability and alleged success.” (Oakeshott) For the first time in intellectual history, Rationalism has gained the upper hand and after four centuries of struggle the European intellect is now fighting for its life. Hence the urgency and apparent stridency in my tone.

Let me begin with the myth of Icarus. We remember Icarus himself, the son, the romantic rebel and his ultimate undoing, but we forget Daddy Daedalus, the ultimate Rationalist and tragic villain of the piece. Daedalus was an engineer, but also an accidental technocrat, designer of the Cretan Labyrinth, hired for his ingenuity as an agent of the state to deploy his devious circuitry to keep its most shameful secret secret. And, as Ovid tells us, Daedalus was very nearly a prisoner of his own device. The rigorous logic of the labyrinth he constructed was impregnable—unfathomable by any man. Yet while the personal and human did not figure in Daedalus’s deadly, impersonal calculations, he had not yet lost the thread of common sense that enabled him to find his own way out, the very clew he would later give to a young woman in love. And so it was, that all his tortuous, logical engineering ingenuity, unrivalled, was soon unravelled.

When King Minos, his erstwhile employer, banged him up in a tower with his son Icarus lest this secret to the Labyrinth get out, he built himself and Icarus each a marvellous flying machine, and, as we all recall but Icarus fatefully did not, carefully instructed him not to fly too near the sun. But as always, what mattered most was all that lay outside the twisting, narrow corridors of all our rationally ordered, Byzantine designs—the greater powers of common sense or love or, this time, beauty. For once again, all our instrumental Rationalism proved impotent in the real world. Icarus, utterly entranced by the unutterable *beauty* of the sun, forgot his father’s meticulously uttered, rational directions, and in the heat of the moment the wax melted and all the carefully engineered contrivances came undone as Icarus fell tragically back down to earth, where he belonged.

In ways I’ll try and briefly elucidate, for 400 years human nature has been tragically at war with our own adopted Rationalism: “We have met the enemy and he is us.” Yet there is still hope in the fact that human nature always wins, even though in the Icarus myth it was a Pyrrhic victory.

An engineer, in cybernetician W. T. Powers’s apt formulation, is “someone who learns whatever he needs to learn to get the job done.” And so the Engineer in us will always be at war with the Rationalist in us. Daedalus was the archetypal technocrat, not because he deployed his expertise as an Engineer, but because of his archetypal Rationalism.

Rationalism can be epitomized as the Myth of Disembodied Knowledge, the myth that knowledge principally takes the form of general propositions independent of the context of their application. The Rationalist takes knowledge to consist of truths that are abstract rather than concrete, holding universally rather than ‘merely’ locally and forming a systematic whole. For him, all genuine knowledge is technical knowledge, knowledge of technique. It can be formulated in rules, encoded in the symbols of language and mathematics, expressed in purely theoretical terms, and put in a book. It is certain, admitting of no exceptions, and applies universally to the objects with which it deals, irrespective of differing local practical

circumstances or contexts of inquiry. Items in the world can be identified once and for all, labeled with the appropriate abstractions, and systematically classified according to their kind. Once we know what abstract *kind* of thing something is we can select the appropriate technique for dealing with it, which follows from our theory of how such things work. Hence Rationalism is the basis also of the contemporary, spurious authority of the expert. The map becomes the territory. Procedure is all.

Wax melted? Please refer to the *User Guide*. Too late? Wings fell off? Please refer to our *Terms and Conditions*.

## A condign punishment

As a society, said Oakeshott, “we remain fixed in a vice of Rationalism.” Of all the social prophets who foresaw our present, benighted age of Rationalism-gone-mad, an age of societal managerialism, with bureaucracy inevitably invading ever more aspects of the once private sphere, the most prescient was the German social theorist and political economist Max Weber. Over a century ago he looked with dread to today’s world of “[s]pecialists without spirit, sensualists without heart; [where] this [collective] nullity imagines that it has attained a level of civilization never before achieved,” as Weber mordantly characterized the modern world. In a “polar night of icy darkness,” it would be a world, he said, of “mechanized petrification, embellished with a sort of convulsive self-importance,” and “the iron cage of bureaucracy”—a technically ordered, rigid, ever more dehumanized society—would increasingly trap individuals in Byzantine systems based on rational calculation and control: a unitary and ever more complex, labyrinthine network of explicit and implicit rules to whom all are subjected equally, in aid of ever greater instrumental efficiency and convenience. Weber’s dreaded bureaucratic technocracy would progressively erode individual human freedom and potential and the richness of human life at every turn, without most people noticing until it was too late. But for Weber, the root of the problem was the fallacy of Rationalism itself, and above all the uniquely modern Western addiction to instrumental rationality.

Weber could never have envisaged the technological boon to Rationalist, bureaucratic totalism provided by our digitally connected world. But a later prophet, Norbert Wiener, the father of the information technology revolution, by the late 1940s clearly foresaw the internet age to which, he was convinced, his own inventions would inevitably in time give birth. Over sixty years ago, Wiener warned of the disastrous social and political consequences of such a world unless we stayed awake to its perils and, as a society, took steps to safeguard ourselves, citing particularly the dangers to democracy itself posed by instant mass communication: we’d be creatures of groupthink-on-steroids.

Wiener, a deeply cultured European mind, was horrified by what he called “the perverse American worship of know-how at the expense of know-what,” putting technological means before conscious choice of ends, the cart before the horse, the Gadget before the Good Life. He foresaw a Promethean tragedy, in which ‘a condign punishment awaits a man’s sin of stealing fire from the gods even where his sin is the fruit not of conscious arrogance but of ignorance of the gods and of the world around him,’\* raising the spectre of a creeping, opted-in fascism, more *Brave New World* than *1984*, even within ostensibly democratic societies.

---

\* The 2<sup>nd</sup> Edition (1954) of his 1950 *The Human Use of Human Beings*, has (on pp. 211-212), “It is a dangerous world, in which there is no security, save the somewhat negative one of humility and restrained ambitions. It is a world in which there is a

Like it or not, today we've all become slaves to the online form, the symbol of our times. Thanks to the ever increasing speed and security of information processing and transmission, and ease of mass data storage, traditional company back-office tasks have increasingly been delegated to the consumer, who, in his own quest for expediency (and with little or no choice in the matter) has become an unwitting accessory in creating a world unrecognizably more bureaucratic than the world a generation ago.

### Prisoners of our own device

The "military-industrial complex" of the old 1960s student-activist rhetoric has now come to include, most notoriously and most damagingly, the banks and the databanks, the Goldmans and Googles, Freddie Macs and Facebooks, and might be renamed "the governmental-financial googleplex." However, the great looming problem is not national security vs. individual freedom and privacy, readily addressable in the democratic world, but the more insidious problem of the bureaucratization of our whole way of thinking, and living.

It is not the erosion of privacy in our lives that is worrisome, but the erosion of our *conception* of our lives. Not the hegemony of the googleplex in narrowing the sphere in which we can keep our thoughts and actions private and exempt from external control, but more its *cultural* hegemony in narrowing and controlling the sphere of our thoughts and actions themselves, whether kept private or not. It is not just the violation of our privacy in tracking our online behaviour, capturing our data and drawing inferences from it that is problematic. Even worse in its societal implications is the increasing uselessness of the inferences made. Worst of all, however, is the unquestioning belief in machine-generated, nonsense conclusions, and in the viability of the whole enterprise. It is less a problem of Fake News than Fake Reality; not of Artificial Intelligence but of Fake Intelligence; not machine learning *per se* but our own all-too-human failure to learn and remember that the darn things are, after all, only machines.

Here Norbert Wiener, all those years ago, warned especially of the dangers of groups of people extending their control over others "by means not of machines themselves but through... techniques as narrow and indifferent to human possibility as if they had, in fact, been conceived mechanically." He saw *this* as the greatest threat of the information revolution to human values, freedom and spirituality, to civilization itself, and above all to our very powers of mind, sense of responsibility and most precious of all in his view, the reality of our freedom to choose.

And this is where the shoe pinches. It is here we must be most vigilant, for as the economist Milton Friedman defined bureaucratic thinking in his "Law of Bureaucratic Inertia," it is the conviction that the only way of doing a thing is the way it is currently being done. Mindless bureaucracy, Rationalism in action, artificially redefines every matter in its own image, and on its own terms. It thinks us into the box.

---

condign punishment, not only for him who sins in conscious arrogance, but for him whose sole crime is ignorance of the gods and the world around him.

"If a man with this tragic consciousness of fate approaches, not fire, but another manifestation of original power..., he will do so with fear and trembling. He will not leap in where angels fear to tread, unless he is prepared to accept the punishment of the fallen angels. Neither will he calmly transfer to the machine made in his own image the responsibility for his choice of good and evil, without continuing to accept a full responsibility for that choice."

The latest, internet-enabled forms of quasi-voluntary bureaucracy require our acquiescence in rigid systems accepted as givens, in which, by definition, our system-permitted responses as humans must be restricted to those responses that might be emitted by machines, and in this way our human functions are increasingly being reduced to machine functions. But of course, we humans are *not at all even like* machines, and cannot be comprehended *as* machines, or *by* machines. We are free, autonomous, intelligent creators and designers.

Machines, on the other hand, are—well—*machines*, neither intelligent nor creative. An “intelligent machine” is an oxymoron and always will be; intelligence is precisely what *cannot* be encoded in rules or algorithms. If you do enough stupidly simple things, in parallel, *very, very fast*, you appear clever. Computers are good at what they do because they have the advantage of being rigidly, *reliably* stupid. As we cyberneticians used to say, it’s difficult to make systems foolproof because fools can be so ingenious. Data is inherently stupid—no bad thing—and as we’d expect, so-called Big Data is *hugely* stupid. The mischief arises when we defer to the data and reify it. And it is a short step from reifying the data to deifying the data.

The data are creatures of what we have previously chosen to measure. And we measure not what it is important or useful or relevant in the context to measure, but what it is easy to measure, or what we already know how to measure, or what everyone else is measuring, or what someone is asking us to measure, and it is a tragic law of nature that *we get what we measure*.

The system *is* what it actually *does*, not what it was or ought to have been *designed* to be. And in this sense the system is always right. If a catastrophe occurs it is *not* the result of an error in the system. Rather, it was a perfectly well-functioning system suited to reliably producing catastrophes like this one. The global financial system did not break down; it did what it was a system for reliably doing, though no one realized what that *was* until the catastrophe occurred.

We soon discover that a system that is always right is always wrong in its deployment in the real world. A machine or bureaucrat cannot compute the value of a variable stipulated as “none of the above.” But in the dropdown menu of real human life, the truth is almost always “none of the above.” And yet in the social and behavioural sciences, from the category-defying of Kinsey to the category-deifying of McKinsey, the mills of academia and commercial pseudo-academia grind out, with ever more refined granularity, their atomized numbers run on artificial, generic abstractions—imponderables made impenetrable. They play arcane, technical, garbage-in-garbage-out tunes on arbitrary invented categories, filling the journals and shelves with their boxes: the fast-moving consumer goods of Rationalist twaddle—a never-ending glut of grant-processed, factory-farmed ‘facts’ which, in the specifically human domain, cut no ice with any self-respecting, organic, locavore empiricist like me...or you.

All their grand theoretical edifices are castles in the air, houses of cards. Yet it is precisely this kind of crude, Big-Science-style, data-driven Rationalism which provides the flimsy foundations for most social, economic, and health policy. All the clever cards and crude conclusions soon fly absurdly in the face of the far wiser intuitions of those who know the territory intimately—those who have their epistemic feet solidly on the *terra firma* of their own rich, intimate local knowledge of the far less highly abstracted—and mostly unformulable—real facts of the matter. In this way, sadly, John Stuart Mill’s “open marketplace of ideas” has mostly been reduced to a derivatives market.

## Grand Turing: Captcha playing gotcha

How many hours of my life have I wasted trying in vain to convince Google's maddeningly obtuse captcha robot that "I am not a robot"? What have we come to? I, a human being, am being made to prove to a robot (usually unsuccessfully!) that I am not a robot? And how often have my repeated failures in selecting the "correct" captcha images been due to my misunderstanding the relevant Americanisms marking parochial distinctions that don't strictly map onto our own English terms, let alone translate directly into any other language?

Of course the tacit assumption is that the whole world speaks American now. It is like the remark overheard by the Russian nurse and writer, Sofya Z. Fedorchenko, during the First World War, made by a Russian soldier in conversation with a German, insisting that Russian was self-evidently the most accurate language in the world: "Take, for example, *noz*; in German it's *Messer*, in French, *couteau*, and in English, *knife*; but *in reality* it is, after all, a *noz*, which makes the Russian word clearly the most correct."

"Select all the images with chairs." As the philosopher Sholom Glouberman reminds us, "*chair*" may have as its nearest French equivalent, "*chaise*," while "*stuhl*" may be the nearest in German, but while "*chair*" can be used for chairs but not for stools or for *chaises longues*, "*chaise*" covers all chairs and all *chaises* but not stools, while "*stuhl*" covers all stools and chairs, as well as some, but not all, *chaises*. But the very same problem arises for native English speakers unfamiliar with the American idiom, or insufficiently fluent in it.

The captcha robocop commands: "Select all the images with street signs." What on earth are "street signs"? It is some Americanism, apparently. Are these what we, speaking the Queen's English, would instead call "road signs" or "traffic signs"? Or are "street signs" meant to be only the ones simply bearing the name of a *particular* street, such as "Kings Road SW3"? Do Americans count "No Parking" signs as "street signs," or not? I haven't the faintest idea. "Select all the images with storefronts." What is a "storefront"? Does a doctor's surgery with a simple brass shingle hung outside, or a funeral parlour, or an elegant basement-level restaurant with only a discreet entrance door a few steps down from the pavement have a "storefront"? Which reminds me: "Select all the images with restaurants." And what *is* a "restaurant" in stateside *patois*? Is an American diner or coffee house or McDonalds or takeaway or ice cream parlour a "restaurant"? It sounds wrong, but you never know; the Americans seem to call all manner of unlikely eating establishments "restaurants," ones we would never dream of calling "restaurants" on this side of the pond, just as they seem to refer to almost any kind of retail establishment as a "store." Nor can I use Google to look up the meaning of some puzzling American idiom, for it's Google's own robocop barring my way in the first place, reminding me of my wrap rage at a new pair of scissors sealed in a tamper-resistant blister pack.

I wonder, though, how many of my captcha failures and frustrations might have been due to my being an admittedly pedantic philosopher who takes language very seriously indeed, probably more seriously than the software engineers writing code all day? Look more closely at the command, "Select all the images *with* street signs": does "street sign" then include the whole piece of street furniture, including the pole? And does it include those images showing only the sign's pole but not the worded placard? If the term is meant to refer *only* to the placard itself, must the *whole of* the placard fall within the boundaries of the selected image? Indeed, does an image of just the tiniest *corner of* an otherwise suitably captcha-qualifying "street sign" count as an image "*with*" a "street sign"? And what about an abandoned, vandalized car with no wheels by the side of the road, recognizable as having once been a car but now incapable of

driving anywhere—does it count as “a car” or does it not? Or again, is a cartoon of a cat a “cat” or not? These are just some of Google’s semantic captcha mysteries I have never fathomed and probably never will; by the time I do, I expect, the loopy technology will have been phased out.

It is no use my asking the robot *what it means* by what it’s asking me to do. It’s not taking questions. The robotic reply would be, “Look, buddy, I’m the one asking the questions, not you. Just answer the question.” Whoever designed the system just decided arbitrarily what *they* wished to mean by some decontextualized American idiom, and all we poor souls can do is try out one random hypothesis after another. “Whaddya mean, ‘whaddya mean?’?” By definition, in a rationalistic, bureaucratic world, there is no room for nuance. Which is why the bureaucratic mindset, including the silicon chipset of the bureaucrat’s mechanical surrogates, operates only in a surrogate world, not the real world. How can a robot, or a bureaucrat trained to think like a robot, even begin to conceive of the richly nuanced real world? It’s like the story of the submarine pilot who was born on a nuclear submarine and has never once been outside, who, when he deftly brings it into harbour and docks it, using his digital control interface, and emerges from the vessel for the first time in his life, blinking into the sun, is congratulated by the harbour master on docking his submarine so masterfully. Uncomprehending, he replies, “What’s a submarine?”\*

I have laboured this point about the captcha robot, at the risk of straining your patience, merely to take it as a familiar illustration of an increasingly ubiquitous phenomenon, highlighting the inevitable mismatch between rationalistic, machine “intelligence” and the real world. As the more astute amongst my colleagues working in AI have been pleading for decades, alas, we will likely never succeed in realizing the dream of Artificial Intelligence. But the booby prize is already here, everywhere around us: Artificial Stupidity.

This particular genus of incompetence applies more broadly than just to computers of course. Sadly, it applies more properly, and even more insidiously, to *people* who have learned to “think” like computers. We clearly cannot blame the computers themselves for behaving so ignorantly. It’s like dogs and their owners—it is often the owners who are badly behaved, not the dogs. The problem is not the arguably necessary obtuseness of computers but the undeniable obtuseness of human beings who, in their failure to understand the limitations of computer technology, and mired in their rationalist epistemology, rely blindly on badly conceived software to do what no software can ever do for them. They have let it off the leash and stoutly turn a blind eye while it does its business.

In the bewildering new commercial landscape, technologically benighted top management remain insufficiently well equipped to properly assess the extravagant claims made by IT sales people talking their own book, and of course the enthusiasts for the latest gotta-have-it technology genuinely believe the gospel they are preaching. What’s more, the executives are daily assailed from all sides by tech-journalistic hype to match. In a hall of mirrors, who can be blamed for losing their way? An august financial publication’s weekend magazine assures us in a feature article that we won’t need doctors’ surgeries any more; in the future we can all be diagnosed and treated online within just a few minutes. What is wrong with this picture?

As more and more executives and boards get taken in by persuasive sales pitches from huge enterprises sincerely hyping the benefits of Big Data and machine intelligence and whatever

---

\* I owe this story to Dr Andrew Bass

next, while warning of the perils of not investing in these prohibitively costly adventures, more and more investment starts to entail more and more effort by management at all levels to defend the massive investment made by a few people in the heat of a “what-were-we-thinking?” moment. At that level, all too often, reality is whatever your peers in other companies are doing. No one dares to be out of fashion because no one wants to be left behind. If the new technology is becoming, by consensus, *de rigueur*, who wants to buck the trend and then have to do the explaining to the shareholders should it turn out the pundits were right after all? The old saying, “No one was ever fired for buying IBM” can sometimes nowadays become, “Someone might be fired for *not* buying IBM,” whenever some alleged technological advance *du jour* becomes what *le Tout-Paris du business* is buying this year, and of course we must have the best!

Over time, the executives have to invent more and more excuses for their pig-in-a-poke, white-elephant technology, until finally finding themselves having to double-think their way around its egregious failings, as its shortcomings and relative uselessness become more evident and undeniable by the day. Well, never mind, it may have proved a disaster, but at least they can still say, to cover themselves, “Who was to know?” But until that moment of rude awakening, the Sunk Cost Fallacy and Confirmation Bias keep mucking in to ensure everyone continues denying the stark reality of the silicon emperor’s nakedness.

Playing the role of the little boy in the crowd in the fairy tale today is a frustrated London supermarket branch manager, assailed by customers wanting their sourdough bread back again. Every day she could sell scores of the fresh-baked sourdough loaves, which fly off the shelves before they have cooled. But she cannot get the stock any more—at most four loaves a day—while having to give to charity racks and racks full of unsold varieties of bread no one wants, and as everyone knows, all the profit from fresh bakery goods comes from the sale of the last loaves of the day. She is told by higher management she cannot have the sourdough because the computer, analyzing extensive data from across the country, has estimated her branch will sell exactly 0 sourdough per day. When she presents the compelling, countervailing facts, fresh from the real world, and pleads with her Marie Antoinette of a manager that her customers have no bread, our intrepid branch manager is told, in an admonishing tone, that she can hardly expect to argue *with the computer!* Can *she* match that much hard data? “Let them eat code.”

The store manager says she can tell any number of stories to match or beat this one. She could increase the branch’s profits if only left to manage her own stock control using her own knowledge of customers and daily observations, left free to override the system when it errs. Instead she’s forced to ignore the evidence of her own eyes and the complaints of her customers, and defer to “idiotic” computers. When customers can’t get their sourdough they go to the supermarket down the road and shop there. After two or three times, they don’t come back. Her staff agree: the corporate fetish for the new techno-wizardry is costing them sales and driving customers away. The staff can’t see how the supermarket chain can make any money at all on this basis, resolutely shutting their eyes to reality. But this is not a cautionary tale from some failing company. This is a thriving branch of one of the world’s most successful supermarket chains. Someone in the executive offices approved the eye-watering spend for rolling out the technology. So of course our beleaguered store manager is talking to a wall.

And here we have it: the triumph over solid empirical reality of data-driven, evidence-based bunk.



## The Real Singularity may be near

The American inventor and controversial futurist Mr Ray Kurzweil and others (above all the distinguished Cambridge statistician, the late Dr Jack Good, who originated the notion) have for some years touted the prospect of a technological “singularity”: a runaway acceleration of progress in technology to the point where an artificial, self-upgrading, networked, super-intelligence surpasses the collective intelligence of human beings and, to make a long story short, takes over the agenda. Some, like Mr Kurzweil, view such a prospect optimistically. However, if such a singularity is truly getting any nearer (which I very much doubt), it is not because machines are getting any smarter—they certainly don’t seem to be, quite the contrary!—but because humans are getting ever stupider in their reliance upon computers to do things that computers cannot conceivably do and never will, as has been known by computer scientists for nearly half a century. Computers may be getting dumber, but humans are beating them in the race away from intelligence. One can all too easily imagine this accelerating decline in human intelligence going into runaway, if it hasn’t already, until we are collectively no smarter than all that mindless IT kit we’ve already built, bought and bought into.

In any event, the good Dr Good’s and Panglossian Mr Kurzweil’s own rather far off and far-fetched singularity is not the one that interests me in the present context. Rather, I’m interested in what I would like to call, by contrast, “the Real Singularity.” The Real Singularity will occur at the point of no return once *The Law of Technological Folly* really starts to bite back in earnest: “*The faster any technology advances and the greater our collective investment in it, the further our dependence on that technology will outstrip our awareness of its limitations and our understanding of how and where to deploy it productively.*”

The Real Singularity will come when nothing works any longer, to all intents and purposes—when information technology stops functioning and grinds to a halt, crashing to earth like Icarus’s own machine, taking us all with it, and simply because we have rationalistically engineered ourselves into a corner, making ourselves dependent on technology that is simply *not up to the job we have assigned to it and never will be*. At that point, the simple things that used to work no longer do, or the things that *do* work take longer and longer to do, far longer than they used to, or simply fail to get done, or people walk, and the whole Heath-Robinson/Rube-Goldberg system starts to collapse under its own weight. That’s the Real Singularity that I think may be coming just around the corner.

I suspect we needn’t even waste our breath in raging against the machine, for in the accelerating frenzy of its *danse macabre*, the infuriating machine already seems to be furiously grinding itself to a halt, all on its own. Perhaps the robots themselves have turned into self-destructive Luddites, like the intelligent bomb at the end of the cult sci-fi movie classic *Dark Star*, stuck in the spaceship’s bomb bay, which self-destructs in a manic frenzy of solipsistic omnipotence.

What I’m calling “the *Real Singularity*,” *if we don’t push back*, may well be near. Some days I suspect it may already be here. But when it comes, at that point we can once more have high hopes for the future, finally beating a retreat from our superannuated blind reliance on technology and systems, policies and procedures, processes and protocols, meaningless bureaucratic rules and regulations, and all the rest of the barcoded barbed wire, and return at last to reality and freedom and the sovereignty of common sense. A great computer scientist once defined “technology” as “shit that doesn’t work yet”; the Real Singularity will be the hypothetical point in the near future when technology is redefined as “shit that doesn’t work any more.” Are we there yet?

Supposedly intelligent, well educated people have increasingly been asking pieces of metal and silicon to make decisions for them. I'm not kidding. There are people (or at least, banks) who really do this. Needless to say, this is not using information technology appropriately for decision support. Rather, this is shamefully abdicating managerial responsibility by impossibly delegating decision making to machines.

My own Real Singularity, again unlike Mr Kurzweil's better-known one, is no longer science fiction. To take a notorious case in point, on a global scale: *We deux-milles-huitards*, by which I mean those of us who lived through the darkest days of the 2008 financial crisis at close hand, will recall the catastrophic consequences of the endemic overreliance upon the Black-Scholes formula for option-pricing and perhaps, most absurdly and destructively of all, the ignorant, blind delegation of decision making, by those old enough to know better, to arcane algorithms they did not even pretend to understand. The actuary David X. Li's ingenious but widely misinterpreted and misapplied Gaussian copula models for pricing CDOs—models whose limitations Mr Li himself understood only too well but those playing fast-and-loose with them and over-relying on them did not—were silently crunching away in the background on uncomprehending machines spitting out “reassuring” numbers, data which senior bankers casually and foolishly deployed by default to go proxy for real decisions, leading them to stay a disastrous course even when a feeling in their water told them all was not well. They were whistling in the dark, and the computer was calling the tune. Don't blame a theory if you have ignored its scope conditions. No model applies globally. Applying Li's models intelligently, and as he intended, would have required taking cognizance of local conditions, while applying a good deal of nuanced interpretation, understanding, intuition and judgement; models rooted in theory, and theory interpreted in the light of unique sets of circumstances. No machine, no matter how “intelligent,” can ever do this on our behalf, or anything like it—let alone ever make decisions for us.

For this is an impossibility-in-principle, not a merely empirical impossibility—as we have known, or should have known, since the discovery of Hume's Law three centuries ago, and certainly since Professor Richard Taylor's eventual decisive triumph over Wiener and Rosenblueth in their high-profile 1943-1950 debate across the pages of *Philosophy of Science*, let alone after the philosopher Professor Hubert Dreyfus's authoritative *magnum opus* in 1963, *What Computers Can't Do*, and certainly after decades of conclusive philosophical and scientific work before and since—not least the cybernetic work of Dr D. J. Stewart and colleagues at Brunel—all leading to the same final conclusion. Despite the wishful thinking of ambitious but blinkered Rationalists in search of their technological Holy Grail, Hume's case rests, and the matter is now closed. Information cannot tell us what we ought to do. Nor need it do so. We human beings can happily do that bit for ourselves—and fortunately, since we are and always will be the *only* ones who can!

Calculating machinery can only calculate, using the algorithms and data we have provided, along with all our assumptions—together with all their inevitable flaws. Like a calculator does when we draw up our household budget, information technology can at most help us estimate what *might* be, based on certain assumptions. It cannot evaluate or decide. It can only deal with abstract possibilities and not the real world, and can merely help us quantify or visualize selected aspects of things in very general terms. It can at best provide us with clues, helping us know what to look for.

But the abstractions of Rationalism, once its great strength, were also its fatal weakness, as Daedalus learned to his cost. The information technology revolution was able to turn both Rationalism's strength and its weakness into a devastating, and apparently ineradicable cultural and intellectual plague.

### Intelligence imperiled

Bear in mind, Silicon Valley sprang up in ideally fertile, Rationalist soil. For Rationalism had proved most virulent in the New World, where it had quickly taken a new and sinister turn, and made itself apparently invincible. As observers have long recognized—from de Tocqueville nearly two centuries ago to Oakeshott in our own time—its inhabitants outdid even the French in their unquestioning Rationalist ideology. The New World's new-and-improved *positivist* strain of Rationalism, leading to deadly, ischaemic *hardening of the categories*, pressed reified, pseudo-scientific abstractions into the service of an all-consuming bureaucratic mindset, nowadays increasingly strangulating the very heart of the Western intellectual tradition, starving it of its lifeblood: free inquiry and reflective thought. This dogmatic Rationalist petrification of thought, knowledge, and understanding remains the most disturbing feature of the bureaucratization of the world. Sometimes inadvertently, sometimes cynically, the counterfeit language of "Science" was now repurposed to the *positivist*, bureaucratic apotheosis of Rationalism in its once uniquely North American variant. Nowadays, however, and increasingly, the Old World does all it can to ape the New, not least in higher education.

The current journalistic fad for Science With a Capital *S* is also founded on the naïve, positivist, Rationalist fantasy that the narrow, quantitative methods of the natural sciences provide the only reliable instrument for revealing the workings of nature, and of humankind. Only during this relatively fleeting moment of contemporary history has the idolization of positivist, Rationalist, organized science become the latest form of organized religion, complete with its own high priests touting their wares in airport bookstalls and televised documentary diatribes. Fortunately, thanks in part to the spectacular public failure of Economics in 2007-08, as it was widely viewed, this dogma so alien to the spirit of real science is beginning to show signs of being on the way out. And nowadays, every educated, cultured person still knows there are numerous equally valid and potent ways besides science for laying bare the nature of reality: art, literature, poetry, religion, music, dance and philosophy for starters. And, though less widely known, there are far more powerful epistemological alternatives to Rationalism and positivism particularly in the sciences—my own focus of work. But there is no time to lose.

For to borrow the words of R. G. Collingwood, "The fate of European science and European civilisation is at stake. The gravity of the peril lies especially in the fact that so few recognize any peril to exist. When Rome was in danger, it was the cackling of the sacred geese that saved the Capitol. I am only a professorial goose, consecrated with a cap and gown and fed at a college table; but cackling is my job and cackle I will." Today we urgently need the cackling of all of us in Europe and beyond, and perhaps most of all in America—all of us who recognize the gravity of the Rationalist peril—those working in the sciences and the arts, and in every field of academia and industry, before much that we all value is lost forever. The problem is a global one, and so must the solution be.

The hopeful return to Reason and to Reality requires us to challenge the insidious error of Rationalism on every conceivable front. In our pressing endeavour to make the world once again safe for the exercise of intellect in aid of human flourishing we must all work together, and

each in our own way, to escape from our 400-year-old cage and take the world of ideas into the future, before a universal, technologically enhanced and entrenched form of positivist Rationalism locks us all in forever.

## The Way Forward

“So much,” as J. L. Austin once said, “for the cackle.” So far, we have merely, and only very roughly, scoped out the problem, albeit one which perhaps poses the greatest threat civilization has yet faced. But what is to be done? And where should we even look for a possible way forward, once it’s time to stop cackling and get cracking?

The dichotomy, let us bear in mind, is neither one between Rationalism (as we have defined it above) and Empiricism, nor between Rationalism and Romanticism, let alone between Rationalism and Irrationalism in any form. After all, can we not still intelligently deploy abstraction in our reasoning without falling into the error of Rationalism and succumbing to the Myth of Disembodied Knowledge? Can we not generalize without trivializing, reason without rationalizing? Of course we can. Can we not use human intelligence to deploy so-called machine intelligence intelligently and constructively? Of course we can, provided we proceed with scepticism and with our eyes open, and our very survival depends upon staying awake. After all, we cannot uninvent the computer and the internet, nor would that help much. Can Reason defeat Rationalism in the end? It can, and indeed it *must* do so if civilization worthy of the name is to stand a chance, if “the outlook for intelligence” (in Valéry’s phrase), is to be anything but bleak.

But can we do *more* than just rant and rave at each opportune moment, and “rage against the machine”? We can and we must; ... though, all the same, we can still push back against the dropdown menu and do all we can, as individuals, to try and break out of voicemail jail, to the extent that we can do so without driving ourselves mad. We can still, and with great justification, rant and rave and rage against the machine, and as Kent McClard rightly intoned in the passage I cited at the beginning, we need urgently to commit, agitate, educate, speak, act, learn, and perhaps especially, disobey; radically alter the way we conduct our lives, change the way we relate to each other, rearrange our goals and desires, and make a revolution within ourselves.

Indeed, here we cannot be complacent, nor set our sights too low, for the task is far more urgent and sweeping today than when McClard wrote three decades ago, let alone when Wiener wrote in 1949, not to mention Weber in 1905. But the task need not be as daunting as it sounds. The heavy lifting has already been done for us. There are firm foundations on which we can build, already laid, foundations which provide in themselves a solution to our dilemma. And unsurprisingly perhaps, the solution, like the problem itself, is to be found in our epistemology. If we are to succeed in making “a risky and hopeful return” to Reality, and “to a position before discouragement and defeat had gained the upper hand,” this revolution within ourselves will be tantamount to no less than an *epistemological* revolution. To this—and exactly what I mean by this—we shall now turn. ...

*...To be continued in Part II ...*

---

<sup>i</sup> There are so many people who have contributed directly and indirectly to the genesis and long evolution of this essay that it feels to me like something of a collective effort—hence this lengthy mini-essay of an endnote to give credit where it is due. The subtitle, incidentally, purposefully tips my hat respectfully both to the psychologist Dr Barbara Held’s *Back to Reality*, (1995) and to the late Professor Stephen Toulmin’s *Return to Reason* (2001). An earlier version of *Part I* of this essay (“A Statement of the Problem”) was prepared as an invited address delivered on 12<sup>th</sup> August 2013 at Mellin Hall, Aalto University, Helsinki, before an audience of some 700 guests paying tribute to the philosopher Professor Esa Saarinen on the occasion of his 60<sup>th</sup> birthday, and marking the launch of *Elämän Filosofii*, ed. F. Martela, L. Jarvilehto, P. Kentta, and J. Korhonen (Helsinki: Aalto University Press / Crossover, 2013), a *Festschrift* in his honour. But the idea for the essay was first sparked by a discussion at the Royal Academy in early 2013 with the German painter-printmaker Massimo Danielis who, in elucidating his 2011 suite of 5 copper etchings, *Icarus II*, stressed that in many interpretations of the Icarus myth, and especially in the popular imagination, there is too much of a focus on Icarus at the expense of Daedalus. Since then I have only been able to read the Icarus myth as a cautionary tale whose moral is not so much how Icarus erred, as how Daedalus erred. My reflections on that brief discussion with Mr Danielis were developed at length in discussions with one or two of my fellow members of our transatlantic “Coffeehouse Circle” (Andy Bass, Beatrice Gutmann, Peter Kenttä, Tuuli Lehti, Isabel Sheinman and Hanna Virta) as it then was—and many of the more important themes and details were inspired by one long discussion in particular with Ms Virta in Farringdon, London at the end of May 2013. My overall thinking on the topic of Rationalism and the themes of this essay was inspired and continually enriched by periodic discussions of the work of Michael Oakeshott with my dear friend Mr Napier Collyns, Fellow of Green-Templeton College, Oxford over many years (beginning in 1991 when I first met him), and it is to Napier that this essay is dedicated. My thinking on the specific topics addressed in this essay has also benefitted greatly from philosophical discussions in the coffee room over the years with Professor Basil Kouvaritakis, whose comments on an earlier version have greatly improved this one, though perhaps not as much as he might have wished. My thanks to Prof. Kouvaritakis and to Mr Pertti Korhonen, Mr Napier Collyns, Professor Raimo Hämmäläinen, Mr Jaakko Korhonen, Mr Jack Martin Leith, Mr Geoff Llewellyn and Mr Jukka Male for their comments on earlier versions, and I daresay I should never have embarked on the present, expanded and revised version had it not been for the constant encouragement of my Editor, Ms Elizabeth Consalvi. Like all of my work, the overall spirit of the piece has benefitted from the personal inspiration of Professor Esa Saarinen, but also, from the shining example, to whom I could never hold a candle, of my teacher of many years ago, Dr Tim Horder, in recent years leading the charge and lighting the way with his brilliant argumentation and ceaseless stream of elegant writings, pulling no punches in fighting the good fight against anti-democratic, anti-educational, mindless bureaucracy, *technopoly* and humbug in the academic sphere. Since completing this essay, I chanced upon the one book by the late Professor Neil Postman that I had not known about or read, his spectacularly prescient 1992 *Technopoly: The Surrender of Culture to Technology*. Though Postman has been a lifelong influence on my work, had I read *Technopoly* before writing this essay I should never have bothered or dared to write it, as he says it all far better than I could ever have done. Should the ideas expressed here resonate with the reader, I would commend his book—long overdue for a reprinting—as essential further reading. It is hard to believe Postman wrote *Technopoly* over a quarter century ago, for every word rings true today, to an even greater extent than he feared, and in ancient times he would justly have been hailed as a prophet. In any case, while I could never have written this little piece without the many friends and colleagues who have played a part in its development, I alone am responsible for any remaining errors or infelicities in expression, and of course for the opinions expressed.