DECISION CHAOS

Unreasonable

Risk management and the attempt that was never made

Kenneth R. Ledger
Decision Chaos

The option that is never assessed

> *The reasonable man adapts himself to the world: the unreasonable one persists in trying to adapt the world to himself. Therefore all progress depends on the unreasonable man.*
>
> *— George Bernard Shaw, Maxims for Revolutionists, 1903*

A risk assessment examines the risk created by a proposed activity. It does not examine the risk created by preventing an activity. There is no line on any register for the harm that the forgone option would have averted, and no line for the knowledge that even a failed attempt would have produced.

This is not an oversight in any particular assessment. It is built into the method. The method takes a proposal as its subject and asks the pessimistic question what could go wrong and how bad could it be. Inaction is not a proposal, so inaction is never the subject, and inaction is never assessed. The consequence is that the discipline is not neutral between proceeding and stopping. It is systematically biased towards stopping, and it presents that bias as prudence.

The bias would matter less if the cost of stopping were visible. In the moment it is not. When an activity proceeds and goes wrong, there is plenty to discuss and review. When an activity simply never proceeds or is prevented, nothing happens and no one knows what was lost. The cost takes the form of an absence, and absences do not appear.

Shaw's reasonable man is not a fool. He optimises within the world as he finds it, his information is good and he is often right when the majority believe him to be wrong. The unreasonable man commits to a world that does not yet exist, his information is poor and he makes progress through many failed attempts. Progress is nonetheless produced by the small fraction of unreasonable attempts that succeed, and there is no method of identifying that fraction in advance, while the failures reinforce the critics who would stop further attempts. A discipline optimised to reduce the number of failures will, by the same operation, reduce the number of successes it could not have predicted.

This paper argues that risk management has become the instrument through which that reduction is achieved, that the resulting cost is deferred beyond the horizon of any process capable of attributing it, and that the correction is available within risk practice rather than outside it.

Two responses to a real hazard

The argument is not that risks are imagined. Almost every case worth examining involves a risk that was entirely real. What varies is the response.

Electrocution in the 1880s was not a hypothetical. The campaign against alternating current distribution made vivid and accurate use of that fact, staged

public demonstrations of it, and pressed for legal limits that would have confined the technology to the safer direct current system. The unreasonable position won out because it could transmit over distance. The danger was real and it was then engineered out rather than legislated away.

Nuclear generation shows the same structure with the ledger visible. Restrictive regulation prevented some nuclear risks, and there is no counterfactual world available for inspection. There are, however, estimates of the other side. Kharecha and Hansen calculate that global nuclear generation has prevented an average of 1.84 million air-pollution-related deaths and 64 gigatonnes of carbon dioxide equivalent emissions that fossil fuel combustion would otherwise have produced. The estimate was challenged in the same journal and is properly treated as contested. Contested is sufficient for the argument, because the figure is not the point.

The Locomotives Act 1865 adds a further complication. It required a person to walk ahead of a self-propelled road vehicle carrying a red flag, and it capped the speed of cars to a walking pace. The hazard was real, but the pressure for the legislation came substantially from established interests in horse transport and railways whose investments the technology threatened. Safety language is not always about safety, and it is difficult to challenge precisely because of what it claims to be about.

Across all three the hazard was genuine and the choice was the same one. Prohibition removes the activity and with it the only source of the knowledge that would have made the activity safe. Engineering accepts the activity and sets about reducing the harm, and it works because deployment is what generates the information the controls are built from. The two are not degrees of caution along a single scale. They are opposite decisions, and only one of them produces anything to learn from.

What the experts said

Between 1519 and 1522 the Magellan expedition circumnavigated the globe, at a cost of four of five vessels and most of the crew, and Elcano rather than Magellan brought the survivors home. Despite the cost it proved many experts wrong and changed global navigation forever. In the 1990s Shuji Nakamura produced an efficient blue light-emitting diode after the field had broadly concluded that gallium nitride was the wrong material and after his own employer instructed him to stop. In October 2024 a returning first-stage booster was caught in mid-air by the tower that launched it, something thought to be dangerous and impractical. In each case informed opinion held the objective to be impossible, uneconomic or unacceptably dangerous, and in each case that opinion was reasonable on the information then available. The experts were not wrong about the risk. They were wrong about the outcome, and those are different judgements.

Ahead of his time

There is a phrase that appears whenever execution of an idea is delayed, and it does a great deal of quiet work. The originator is said to have been ahead of his time. Tsiolkovsky published the rocket equation in 1903 and nothing built on it reached orbit for fifty four years. The phrase absolves everybody involved, because it locates the cause in the era rather than in any decision.

It conceals three different delays. The first is a genuine dependency. The idea existed but a required input did not, and no decision by anybody would have brought the realisation forward. Tsiolkovsky is largely of this kind. Orbital flight waited on metallurgy, turbopumps, guidance and a quantity of capital no private party was going to supply, and it is difficult to identify a refusal that materially delayed it.

The second is a foreclosure. Every required input was available and the delay consisted entirely of the idea not being permitted. Wegener proposed continental drift in 1912 and the profession dismissed it for half a century, not for want of evidence but because the proposal did not fit. Marshall and Warren identified the bacterial cause of peptic ulcers in 1982 and were rejected long enough that Marshall eventually infected himself to make the point. Nothing was missing in either case except agreement.

The third involves no refusal at all. The problem is solved, the solution is correct, and the interval between solving it and being able to sell it exceeds the horizon of any process that allocates capital. Kinoshita demonstrated extreme ultraviolet lithography in 1986 to a deeply sceptical audience, and the technology entered high-volume production in 2019, thirty three years later. Nikon abandoned the effort in 2011 as costs escalated, and the single remaining firm, ASML, survived only because its customers funded it directly from 2012. The arrangement had to be invented, because ordinary capital allocation could not carry a thirty three year position. Every advanced semiconductor now depends on a technology no conventional appraisal would have funded to completion.

The three are indistinguishable in retrospect and the phrase collapses them. A test is available. Name the input that was absent. If a specific material, instrument, technique or quantity of capital can be identified, the delay was a dependency and nobody is at fault. If no such input can be named, the delay was a decision, and the fact that it was collective and well-intentioned does not alter what it cost. If the absent input was patience, the delay was structural and will recur on the next problem of comparable duration.

The credit the solver does not receive

A further consequence concerns the person rather than the technology. Whoever solves a problem of this kind will not be the person who benefits from it. Tsiolkovsky died in 1935, twenty two years before anything reached orbit. Kinoshita is a

university professor and holds no share of the industry his demonstration made possible. Nakamura received a bonus of twenty thousand yen and had to litigate against his employer to obtain anything further.

This is not a complaint about fairness. The standard test applied to any significant commitment is whether the party bearing the cost will capture the value. Foundational work fails that test by construction rather than by weakness, because the value of a foundation is captured downstream, by other parties, after the people who took the risk have gone. An appraisal that asks who captures the value is structurally inclined to reject foundational work, while appearing entirely rigorous.

It also means the behaviour cannot be induced. A person who commits to a problem whose return cannot accrue to him is not responding to an incentive, and no incentive can be designed that would reproduce him, for he is unreasonable. The only variable available to institutions is whether he is permitted to proceed, which is precisely the variable that risk and appraisal processes exist to control.

History, time and the misallocation of credit

Innovation is usually dated by its arrival rather than its origination. The system in place when an idea finally succeeds therefore tends to receive credit for work whose foundations may have been laid decades earlier.

Tsiolkovsky published the rocket equation in 1903 under Imperial Russia. By the time Sputnik flew in 1957, the political system had changed completely. The Soviet system contributed direction and execution, but not the origin of the chain it inherited. Germany shows the same lag in a different form. Much of the scientific foundation associated with its later achievements had been laid before 1933, while the regime that followed removed or drove out many of the people who had created it.

The Cold War produced genuine new foundations of its own, including packet switching, public key cryptography, molecular biology and advanced aerospace technologies. The point is not that one era innovates and another merely consumes. It is that every era inherits ideas laid down earlier while creating some of the foundations that will only become valuable later.

That delay creates a measurement problem. A system can appear productive for years while drawing on an inherited stock of ideas and creating few new ones. The eventual successes are then credited to the period in which they mature rather than the period in which they were made possible. What is rarely measured is whether the current system is replenishing that stock at the same rate it is consuming it.

Mistaking execution for origination

The common feature of these cases is not the politics of the system but the shape of its decision making. A focused decision system becomes very good at selecting

between objectives it can see, comparing them and directing resources towards those judged worthwhile. Reasonable men operate such systems extremely well. The difficulty is that origination does not begin with a set of alternatives waiting to be ranked. The unreasonable man is valuable precisely because he produces possibilities the system did not know to ask for. The result can therefore be strong execution alongside weak origination, a condition that remains invisible for as long as there is an inheritance of earlier ideas to execute against.

Nothing about this requires a state. A company that routes every proposal through a stage-gate, a business case and a risk assessment has built the same filter, and it will produce the same profile: competent delivery of assigned objectives, and a steadily diminishing supply of objectives nobody assigned. The resulting pattern can be the same, even though the coercion is absent, which is what makes it difficult to see and easy to defend.

There is direct evidence that institutional conditions affect the production of new ideas, not merely their execution. Azoulay, Graff Zivin and Manso compared investigators funded by the Howard Hughes Medical Institute (HHMI), which tolerates early failure, rewards long-term success and allows considerable freedom to experiment, with similarly accomplished scientists funded through National Institutes of Health (NIH) project grants. The HHMI investigators produced high-impact papers at a much higher rate and moved more often into novel lines of inquiry, while the total number of publications changed little. Same discipline, comparable calibre, similar resources. What differed was the environment in which ideas were generated.

This matters because the system that ultimately develops or deploys an idea is not necessarily the system that created the conditions from which it emerged. Innovation is usually credited when it becomes real, visible and valuable. By then, the environment that generated it may have disappeared, leaving the inheriting system to receive credit for a capacity it may not itself possess.

The lag then creates a second problem: by the time a decline in origination becomes visible, the environment that caused it may be decades in the past. Nobel Prizes are commonly awarded twenty to thirty years after the work they recognise, so the prize record can flatter the present by describing the past. Gene editing tools rest on bacterial immunity research from the 1980s and 1990s, and the messenger RNA platforms deployed in 2020 rest on a line of research that struggled for funding before producing the pivotal 2005 work. Bloom, Jones, Van Reenen and Webb find research effort rising sharply while research productivity falls, and the Organisation for Economic Cooperation and Development reports multifactor productivity growth remaining weak. Whether this reflects risk aversion, the terms of funding, conservatism in review or the genuine difficulty of later-stage science is unresolved, and this paper does not claim to resolve it. It claims only that nobody is testing it, because no institution is responsible for testing it. That is the severed

feedback loop described in The Deferral, operating on a research system rather than a firm and over a generation rather than a reporting cycle.

The present instance

The progression from the internet through social media and cloud computing to data centres and artificial intelligence has followed a consistent pattern. Benefits are discovered, dependence increases, and attention then shifts heavily towards harms. The introduction of television followed the same arc, and the arguments made then about attention, passivity and social cohesion are recognisable in the arguments made now.

Data centres are the instructive case because the objections are quantifiable and legitimate. Electricity demand, transmission capacity, water consumption, land use and emissions are all real, all measurable and all properly the subject of public decision. Nothing in this paper suggests otherwise.

That is precisely why the choice of response matters. A control-oriented framework asks how the harm can be prevented. An engineering framework asks how the benefit can be obtained while the harm is managed, and accepts that most of the knowledge required to manage the harm will be produced by the deployment itself. The war of the currents is the closer precedent than the red flag, and it is worth asking which of the two the present debate more closely resembles.

What can be done

Every assessment of a proposal should carry an assessment of the option being declined, including the harms the proposal may have averted and the knowledge the attempt would have generated. The assessment should record both, so that a decision to stop is argued rather than assumed. A recommendation not to proceed is a decision with consequences, and it should be documented as one.

A discipline that evaluates decisions by their outcomes will always prefer the choice that produces no perceived negative outcome at all. The evaluation must instead be of the quality of the decision at the time it was made: what was known, what was reasonably knowable, what alternatives were compared and on what basis the choice was made. Under that standard the unreasonable man is not vindicated by success or condemned by failure. He is assessed on whether he understood what he was taking on, which is the only question a risk function is equipped to answer and the one it has largely stopped asking.

Conclusion

Underneath all of this sits a difference in posture towards disorder. Ideas arrive messily. They require play, false starts, wasted effort and long periods in which nothing defensible can be said about them. That is the condition in which the

unreasonable man works, and he does not find it uncomfortable. Risk assessment is built for the opposite condition. It is structured, it seeks to reduce variance and it treats disorder as a defect to be removed. The reasonable man is not hostile to new ideas. He prefers control, and control is the one thing that idea generation cannot survive in quantity.

The two are therefore not personality types but two responses to chaos. One works inside it and accepts that most of what emerges will be worthless, but within that there are game-changers. The other suppresses it outright and accepts that the worthless and the foundational will be removed together, because at the point of decision they are indistinguishable.

Risk management started out with the intent to help people understand what they are taking on. It has drifted into being the mechanism by which they are stopped, and it has done so without anyone deciding that this should be its purpose. The drift is correctable, but only once the profession accepts that a decision not to proceed is a decision, that it carries consequences, and that those consequences belong on the same page as every other risk it records.