

Kaleidoscopic Change¹

James Wilk

University of Oxford and
Interchange Research (London, Helsinki, New York)

A new era is dawning for our field. While it is a cliché to note that the pace of change appears to be accelerating as never before, if only for this reason it is undeniable that management consulting expertise has become more indispensable than at any time in the past. New methodologies will be sought by our clients, and developed and deployed by our profession, to tackle the seemingly unfathomable complexity of all the disruptive developments and disruptive new technologies emerging.

However, I believe that it is our own profession, over the next couple of decades, which will be responsible for delivering one of the most disruptive and transformational developments of all, and one which will inevitably lead to sweeping, disruptive changes across our own industry: *the ability to catalyze major, enduring, desired transformations rapidly, with minimal risk, and with scientific precision.*

If that sounds a tall order, and if nothing can prepare you for what I'm about to say, it is *not* because it is at odds with your experience of the world. Rather, it is at odds with what we've all been taught *to make of* that experience. And that's the crunch.

Reverse butterfly stroke

You know about “the butterfly effect” from Chaos Theory in mathematics? How a butterfly flapping its wings in China can alter the weather over New England? Suppose you could turn it around, and figure out in advance, OK, *which* butterfly, on *which* gatepost in Shenzhen (say) would have to flap its wings at *what* time, at so many beats per second for exactly *how*

¹An earlier version of this piece was delivered as a Keynote Address to *f.ounders*, a private meeting of the world's top 100 technology founders, held in Dublin from 28th to 31st October 2010. My attempt at a brief, oral presentation of extensive, complex material was only coherent at all—and is only readable today, six years later—thanks to the generous editorial assistance, in the run-up to my presentation, of four gifted writers, who also happen to be four of the smartest and most delightful people I've ever had the privilege to know: Herb Addison (ex-Oxford University Press), Beatrice Gutmann of Fordham University, Quentin Hardy of *The New York Times*, and Jaime Wolf of *The New Yorker* and *The Atlantic*. The writer and reader alike owe much to them, but any remaining muddled diction is all mine. **I'll be giving a WCoMC Shop Talk on this topic soon, date to be arranged; will I see you at The Hand and Shears? Or after this paper will I be drinking on my own?**

long, to guarantee a beautiful Labor Day weekend on Cape Cod? What if you could do the equivalent of that in society, or in business, or in the development of new technology or new applications? Or in transforming the culture of an organization at a stroke? And what if it worked, reliably, every time?

What I'm talking about, however, has nothing to do with Chaos Theory—it hinges on a different set of scientific concepts altogether. What I'm talking about is the analysis of highly complex systems in order to *deliberately* create major transformations swiftly and reliably by pinpointing and implementing such irrelevant butterfly wing-flaps.

Here are some real examples from my own work. A major government department was under siege from all sides and on the point of catastrophic collapse, with open conflict amongst the top staff and with other agencies, leading to widespread calls for the removal of its chief executive. That organization was transformed overnight into a harmonious, well-oiled machine . . . simply by the CEO deliberately *moving an ordinary glass coffee pot twelve inches*, from one side of a partition wall to the other.

In another case, a major department central to the operations of a leading media company was struggling to achieve its objectives. Its effectiveness was transformed within days, after a senior executive asked the new CEO an innocent and *entirely irrelevant question about baseball*, knowing in advance that asking that question would swiftly secure the desired transformation.

One more: A company faced the imminent loss of half its revenue, but instead maintained it, prospered and entered a stage of rapid, dramatic, sustained, market-leading growth, owing to the chain of events purposefully triggered by a carefully crafted, 37-word telephone message.

Each of these catalytic actions was precision-designed in advance through the rigorous, scientific analysis of an unimaginably complex system.

Scientific detective work

I am not citing these seemingly innocuous pivots to trivialize my own work. The catalytic, systemic interventions were indeed trivial, and in my experience they *always* are. But the results were massive and immediate, *occurring exactly as predicted*: and that's what this is all about, and what I find most exciting of all.

Many of us here in the Company already know the impact on the world that a few lines of computer code can have—work hard over a weekend, and you have the basis for a commerce-changing eBay, a viral meme, or a WhatsApp. Others are aware of such precision-designed campaigns in halting the spread of disease outbreaks. Large organizations, geopolitical situations, global markets and the growth and transformation of businesses are no different.

What we've been waiting for is a reliable method of analysis to enable unimaginably complex, large systems to be lastingly transformed in this way—by design. Or again,

changes and major advances, especially in the tech world, result from an epiphany. But what if you had a reliable method of analysis to design such epiphanies to order?

Those means are now available. In fact, for *any* desired change in the world of affairs, or for *any* practical problem, no matter how big or complex, it is now possible—systematically—to rapidly identify the smallest action you can take that will be certain to flip the situation from the existing state to the desired state—in the shortest possible time, with fewest resources, with a minimum of risk and with no undesired side effects.

And most of the time, the actions required are as easy and trivial-sounding as the baseball question, the telephone message or moving the coffee pot. Today, in 2016, such a startling capability is no longer science fiction, but science fact.

It has been my life's work. Over the years, my quest for a General Theory of Intervention took me from mathematics, logic, semantics and academic philosophy through neuroscience, psychology, psychoanalysis and psychiatry; through embryology, cell biology and semiotics to, perhaps above all, cybernetics, as I hunted down my elusive quarry.

The scientific methodology for precision-engineering this type of code for predictably transforming a complex system is called "*minimalist intervention*." The technology consists of a very rapid, rigorous process of scientific detective work, methodically posing questions of someone who knows the territory of their situation in intimate, idiosyncratic detail, and who then implements the pinpointed action to secure the transformation desired.

Just to give you a sense of what I mean, take *the baseball question*. The question was, "Which is it going to be, the White Sox or the Cubs?" As a result of the analysis, the executive knew that this innocent question would flip the dire situation overnight. How? After much analysis, it became clear how to quickly create the desired state provided we could enlist the CEO's support—never mind the details—but this still required getting the CEO aside for a few hours to think through the necessary moves with any one of the top execs, and everyone had tacitly assumed, on good grounds, that he'd never in a million years put aside the time, indeed not more than 10 minutes—*especially if he knew the topic*—and that he'd certainly dismiss it all out of hand.

But further filtering for what was required revealed that sitting next to him at a ballgame would enable the right kind of conversation. The exec knew enough about patterns in the CEO's personal idiosyncrasies to realize that *this* particular Sox/Cubs question asked at such and such a time would be sufficient to get himself invited to a game at Wrigley Field with the CEO the next weekend. The CEO was a baseball fan newly relocated to Chicago in advance of his family, with no new team picked out to root for, and no one to go with him to a ballgame; *and he'd just been given two tickets to a Cub game*. Neither the Cubs nor the Sox had *anything* to do with the *problem*, but a ballgame had *everything* to do with the *solution*.

Any *minimalist intervention* depends on idiosyncratic details of the situation quite irrelevant to the problem, for it is impossible to know from a statement of the problem and its consequences just *which* idiosyncratic constraints must be lifted to release the change desired. The solution—a catalyst that releases the desired change—rarely has any relevance to the

problem it solves. You can forget about the problem altogether, in fact, and instead focus your attention on the *desired* state-of-affairs and what it will take to release it.

Findings

You may be thinking, “Big deal—he named a moment and declared victory,” or “this is just Organizational Psychology 101,” or “he got lucky.” The first few times I too asked myself questions like that. The big deal is in X-raying the intractable haystack to find that one simple and obvious needle (more commonsense with hindsight than organizational psychology) and declaring victory upon finding it, knowing in advance that it will do the job. Most minimalist interventions require no psychology at all, but the interventions obvious with hindsight would appear instead as “Software Engineering 101” or “Corporate Finance 101” or what have you. Getting lucky is a matter of chance; getting lucky every time is a matter of science.

In some ways, the idea of minimalist intervention should not be startling. You know how it is when you look back on any hard-won success of yours in your life, your career, or in your business—and how, with hindsight, you can often identify just one small, key thing you did—one lunch, one phone call, one tiny shift in how you were doing or looking at something—where if you’d only known at the outset to do *that one small thing*, it could have saved you 90% of your struggle, and got you there a lot sooner.

The exciting thing is the *technological capability to pinpoint these catalytic actions in advance*. We haven’t just done this four or five times. To date, my scientific colleagues and I have completed the design of some thousands of successful interventions of this kind, including over 750 where the interventions designed instantly resolved major, mission-critical issues—issues valued in the millions, tens or hundreds of millions, or in some cases billions of dollars, and many of these resulted in game-changing, industry-transforming developments.

The upshot is huge—philosophically and scientifically as well as practically—and requires us to shift our whole view of the very fabric of reality. Among the purely *practical* findings from our research:

Change does not take time, and when it occurs, it happens all at once.

The scale and scope of a problem, and the amount of time it has persisted, are irrelevant to what it will take to dissolve it.

Indeed, problems don’t come in sizes; there are no big problems and small problems, only lesser and greater consequences.

Any significant change can be achieved instantly and painlessly if it can be achieved at all.

The desired change can always be found to be already immanent in the existing situation and needs merely to be released by pinpointing a tiny, minimalist intervention into the system.

Such an intervention rarely involves more than a subtle, innocuous, single communication, often delivered to a single individual, or at most, to a few individuals.

The result is then an all-or-none flip from the existing state to the desired state all at once, across the board, with nothing in between.

The most startling thing of all, however, is that the analysis to pinpoint the minimalist intervention that will release the desired change typically takes two to three people about four hours; sometimes two sessions totaling eight hours; almost never as many as twelve hours.

My colleagues and I have found all these principles to hold universally, and in our experience to date *there are no exceptions*. A minimalist intervention can *always* be found. I'm talking about finding baseball questions or moving coffee pots or 37-word phone messages that will predictably create transformations in days or weeks that had been judged by all the experts to take months or years.

I remember the CEO who insisted it was a known fact that a particular desired change could not be achieved until the entire structure of her industry changed, which everyone knew would take at least ten years. That "ten-years-away" change took four hours to design and was complete, for her organization's purposes, within 24 hours of the session. Even changing the corporate culture of an organization can take not years, but weeks, or even *days*, by designing a minimalist intervention.

Kaleidoscopic change

None of these examples of minimalist intervention make sense, in our usual ways of making sense. There are no case histories to be related, no narratives that can show how it all played out. This is a crucial point, saying much about how people have long misunderstood complexity. Wearing our current cultural blinkers, we think in narratives, but the world is not narrative in structure, following plot lines. Reality is infinitely multi-dimensional, and has none of the sequential, narrative structure we impose on it. Rather, reality is more nearly kaleidoscopic, and change occurs kaleidoscopically.

In a well-made kaleidoscope, if you apply the most delicate of disturbances by rotating all or part of it, you observe sudden, all-or-none transformations of colors and patterns out of all proportion to the magnitude of the disturbance applied. The image does not shift by degrees, but flips over from one complex symmetrical pattern to a radically different one. Yet the jump does not take the form of a *random* leap, as if you were watching a slideshow, moving from one slide of a stained glass window to another. The kaleidoscopic shift appears as a natural, unfolding or metamorphosis of the image, as if one were watching, in time-lapse photography, the embryogenesis of a plant from a seed, or the opening of the petals of a rose.

Similarly, when we are seeking to create any desired change, it is the observed *pattern*, and transformations of the observed pattern, in which we are interested. You cannot "slightly" alter a pattern without it then being an entirely new pattern. If you want to change the pattern

that you observe through the eyepiece of the kaleidoscope, there are a number of options at your disposal. The pattern you observe through the eyepiece can shift in a number of ways (try this sometime): You can alter the distance of your eye from the eyepiece; you can change your viewing angle; you can alter the angle between the mirrors, or you can nudge the colored glass shards.

By understanding the points of view of the players and the way in which these points of view produce the patterns of interest to us, as well as the way in which the shards happen to be reflected from these different points of view, you can precision-engineer seemingly trivial manipulations of the “shards” of mundane reality to produce, in breathtaking all-or-none fashion, the precise transformations desired.

The thing is, the necessary shift can be very rapidly identified, in a matter of hours, because we are *filtering* the complexity, not seeking to understand it. What we’re addressing here is, after all, unmappably complex. Fortunately, we do not need to map or model *anything*.

A scientific model answers but one question or class of questions at a time; every time you ask a different question, you have to throw away your model. The model of light as moving in a straight line from a point source explains why your shadow is shorter at noon than at sunrise, but we need a model of light as a wave to explain the refraction of light by a prism. No matter how good our model, if it is answering the wrong question it will be as good as useless. The map of the London Underground is very useful, but not for navigating the city on foot or negotiating the one-way streets.

The point is not to model complexity, but to filter it.

At the outset of the analysis, the existing state and the desired state may appear to be as far apart as can be imagined. That’s a function largely of how they have been described. For reality is infinitely redescrivable, and there are an infinity of descriptive dimensions that may be salient. In the analysis what we’re looking for is the descriptive space in which the existing state and the desired state are already so near to one another, that one will merge into the other with the smallest possible nudge, simply by shifting one or two key constraints.

Finding that space is at the heart of this new methodology. When people describe their problem, they only describe it along an arbitrary set of dimensions, and like the proverbial drunk they look for their lost house key under the streetlamp, not because that’s where they dropped it but because that’s where the light is.

Clearly, this has nothing to do with “nudges”

By now it should be clear that we are absolutely *not* talking about the kinds of “nudges” familiar from behavioural economics, as popularized by Cass Sunstein and Richard Thaler. Far from it. In fact, in many ways, a minimalist intervention is quite the *opposite* of a “nudge”!

First of all, where a so-called “nudge,” as you know, is meant to be *incremental*, insofar as it

aims to incrementally influence people's choices in a desired direction, a minimalist intervention is *transformative*, all-or-none, delivering a permanent, large-scale transformation instantly, across the board.

Second, whereas a “nudge” is meant to be *sweeping* in its implementation, and thus tends to be generic, scalable, and aimed indiscriminately at a multitude of unnamed individuals, a minimalist intervention is *pinpointed*, custom-designed and *single-use*, most commonly targeting no more than one to five named individuals in catalyzing an across-the-board transformation.

Finally, while a “nudge” is *empirical* (“hit or miss”), founded on astute hunches or on research-based trial and error, with a probabilistic outcome at best, a minimalist intervention, by contrast, is precision-engineered and designed to be sure-fire, catalyzing the whole of the desired outcome precisely, with 100% certainty.

This is in no way to disparage the excellent and important work being done on nudges. It is simply a different field of work, based on very different theory.

But while the behavioural economists have, to their great credit, at least dipped their toes into the water of a radical new epistemology, the science behind minimalist intervention dives deep into that vast, unexplored, newly discovered ocean, arriving finally at the sea bed, the very foundations of our human understanding, and bumping our heads against it.

Dare we enter the Information Age?

To make better sense for you of Minimalist Intervention, we need merely to view reality from the perspective of contemporary science. I say, “merely,” but minimalist intervention makes total sense only once we abandon the 400-year-old rationalistic view of the world and substitute a view fitting more closely with the way things really work.

This should not trouble us too much. We do, after all, live in the Information Age. What we do not do well, *yet*, is *act as if* we do. The view of the universe and of causation that almost everyone unquestioningly adopts was concocted by men in tricorne hats and powdered wigs, knee breeches and frock coats. Actually, most of it was concocted by their great-grandfathers in turn, barely out of the Middle Ages, before the dawn of modern science.

They did a good job of describing the world, considering the tools they had to measure and manipulate the world. Today's tools give us a far better sense of complexity, and with that we must deploy a very different understanding of causation—even to the point of abandoning that concept altogether.

Since the advent of cybernetics we have had better tools, and we must meet them with our own, *new* ways of thinking, not ways of thinking dating from before The Fall—of Newton's apple. If we are to understand how the world really works, we need first to abandon the pseudo-scientific fairy-tale of The Atoms and the Molecules: the antiquated 17th-Century narrative of colliding corpuscles, matter-and-energy, cause-and-effect; the myth that the

concepts of a pre-nineteenth-century physics and chemistry provide a sufficient basis for explaining the mechanics of the universe; the fantasy of all behavior as being causally explained in terms of pulsing nerve cells and neurotransmitters, and of living things parsed in terms of crude, chemistry-set conceptions of the role of DNA in biological inheritance.

Likewise, we must jettison the notion that complex human events, from businesses and markets to wars, are in an entirely separate category of analysis.

Truly revolutionary scientific work over the past 85 years, work cutting across disciplines, has reconceived the fundamental mechanics of the universe in terms not of matter and energy, power and forces, cause and effect, but in terms of form and pattern, flux and constraint, information and design. The very fabric of reality has been reconceived by these investigators in terms of infinitely multidimensional networks of information flows, patterns of purely communicational interaction, at every level of description.

Purpose and desire, perception and context can be shown to form a fundamental part of the mechanics of the physical universe itself and can be rigorously compassed in hard scientific terms, dissolving entirely the superseded Galilean divide between human and natural domains, between subjective and objective. If the very fabric of reality consists of information at every level of description, we can derive “*it from bit*” (to borrow Wheeler’s phrase). When we understand this we will have truly entered the Information Age.

How the world really works...and how to change it

In my own scientific work, we begin with the assumption that continuous, random flux is to be expected anywhere in the universe, and therefore it is the *persistence* of any invariance, any pattern, that needs accounting for. And we account for it in terms of the constraints that mean that nothing other than the observed pattern is currently possible, given those constraints in place.

By way of such a figure-ground reversal, putting all scientific explanation in negative terms, we replace the old concept of cause-and-effect with the far more flexible and more powerfully explanatory concept of flux-and-constraint. We take any state-of-affairs we seek to explain, or to change, and ask which constraints are in place such that no other alternatives are currently possible. And we are always asking, implicitly, not “why this?” but a much more specific, “why this *rather than* that?”

If continuous flux is to be expected, and any persisting pattern is the result of the constraints that are currently in place, then change is never caused or brought about; it is only ever released. And we can release the desired change swiftly and with precision if we can pinpoint the idiosyncratic constraints that need to be lifted and/or inserted in some state-of-affairs to flip it over into the desired state-of-affairs.

Desired change can be instantaneous when the right constraints are lifted and inserted at the same time such that where before only the existing state was possible, now, all at once, only the desired state-of-affairs is possible.

On this view, we understand the fabric of reality as consisting of *differences* on the move (information) and *imparities* on the move² (unequal weightings, values, determining which differences are to count as error). There is matter and energy on the move too, to be sure, but no longer in the driver's seat, or rarely so, and the laws of physics are simply among the countless operative constraints we must take into account. The constraints that are the stock-in-trade of the physicist are rarely the most salient constraints in determining what happens in the real world outside her laboratory. The constraint embodied in a well-known law of nature keeps me from getting to Birmingham faster than the speed of light; but the cancelled 5.17 constrains me from getting to my destination before "eight o'clock at the earliest."

Change, on this understanding, is *a difference over time*, and by definition any difference is a descriptive difference. To say that something has changed is to identify it as the same object as it was previously, but altered in some descriptive respect. Any descriptive difference will be relative to the particular aspects we happen to be taking into account—relative, that is, to our point of view.

Like information itself—differences on the move—*change, description* and *point-of-view* are now seen as fundamental to the mechanics of the physical universe, and can be rigorously quantified. We are *not* just saying that what something signifies, in context, for someone or something else, 'is *as important as* what it is', but on the contrary: that what something signifies in context from different points of view *is* what it is.

Envoi

The technology of minimalist intervention has been tried and tested, and used to create thousands of major changes in the business world. Having got ample proof of concept, now it is simply a matter of scaling it.

In 2016, together with my Interchange Research think-tank colleagues, I embarked on a long-term venture to scale our proprietary minimalist intervention technology (in partnership with a selected, global management consulting firm, to be identified), and to train up the next generation of professionally certified Minimalist Intervention practitioners, eventually on a global scale.

Yet even the dramatic, demonstrated success of minimalist intervention in solving previously intractable problems and creating major transformations in impossible time-frames, constitutes merely one upshot of the alternative take on the world I have been sketching above. There are many more such upshots already well on the way, not least in the psychiatric world, for example.

For adopting a different take on the world means we're dealing with a different world—this time, one where the possibilities are infinitely greater. According to the Old Official View, there is on the one hand the way things are—how the world just is. On the other hand there

²as first charted scientifically by Dr D. J. Stewart at Brunel University in his Ternality Theory and Ternary Analysis

are our perceptions and conceptions of how things are, a kind of picture in our heads. We reject this view entirely. It is not at all a matter of territory and map. We are only ever dealing with the territory-*as-mapped*.

What's more, our perceptions of the world are not in our heads at all, but are objectively out there with the rest of the objective world. What we call "our perception of the world" is just *the-world-as-perceived-by-us*; it's the objective world as we correctly or incorrectly take it to be. This has one immediate practical upshot in itself: 'Our situation' is never a given, but is how we currently situate ourselves—and *even if* it can be objectively measured and validated, that still doesn't make it relevant or real. Our situation is how we situate ourselves in relation to what we think needs doing, how we take things to be and why, and what resources we believe are needed and at our disposal. All of this is a highly negotiable matter, never a given.

For example, a leading financial services institution was beset by long-running intractable problems in the area of risk, costing it many billions of dollars. It had just instructed a leading, global, management consulting company to embark on a five-to-seven year consulting assignment with an agreed minimum price tag of over \$50 million dollars, involving over 400 consultants, and headed by the consulting firm's eminent Partner who led its Risk Management and Internal Control consulting practice for a decade. However, when, at the client's behest, I managed to resolve the entire problem with a single minimalist intervention—designed in less than eight hours from start to finish and implemented within a fortnight—the other firm's management consulting contract was cancelled before it could even start as the problem had already been completely solved. This was possible not least because by the end of those eight hours "the situation" was no longer what it had been before, and had long been "proven" and "documented" to be. Situate yourself in one way and it will take seven years and 400 consultants and scores of millions of dollars to solve your problem; situate yourself another way and you can get the same or indeed better results in a fortnight.

This example should also give you a taste of just how disruptive this new technology has been demonstrated to be, and how it will inevitably transform the power of management consulting over the years to come, and enable us to transform our clients' organizations on a scale never previously imagined. Our profession, over time, will be elevated to new heights.

It is a new epistemology that is at the heart of the future transformation of our industry. In place of a unified reality, "how things are," to be parsed according to a unitary privileged narrative such as "the world according to physics," we have instead a more fertile alternative. From an infinity of mutually irrelevant and only coincidentally connected *ways* things are, shifting constantly in response to shifts of point of view, we can *change the question* and thereby change 'how it is'—often radically, sometimes kaleidoscopically.

Dr James Wilk

Faculty of Philosophy, University of Oxford

Managing Partner, Interchange Research, London

Freeman of the Worshipful Company of Management Consultants

Email: james.wilk@seh.ox.ac.uk