

An Uncertain World II: *Adapt*, by Tim Harford

04 June 2011

Contents

1	An Uncertain World II: <i>Adapt</i>, by Tim Harford	1
1.1	The thesis	1
1.2	The contradiction	3
1.3	Consequences	3
1.3.1	Capecchi and Tharp	4
1.3.2	Whole Foods, Timpson, Chile, and Wal-Mart	4
1.4	Centralized experimentation	5
1.5	The Financial Crisis	8
1.6	Foxes and Hedgehogs	9
1.7	If not trial-and-error, then what?	9

1 An Uncertain World II: *Adapt*, by Tim Harford

There is a contradiction at the heart of *Adapt*, the new book by the Financial Times' Undercover Economist, Tim Harford, and once the contradiction is unpicked, the rest of the book unravels.

In this review I hope to state Harford's thesis, demonstrate its internal inconsistency, and then show how *Adapt's* claim to introduce "a new and inspiring approach to solving the most pressing problems in our lives" (dustjacket) reduces to a set of cherry-picked examples in order to advocate, misguidedly, for the extension of free-market ideas to other parts of our lives.

1.1 The thesis

Facing an unpredictable and interconnected world, Tim Harford is certain that we must jettison faith in certainty.

Strangely for a book written after the financial collapse, Harford's thesis is that, to deal with the complexity and unpredictability of the modern world, we should take our inspiration from the market, and apply its methods to other parts of our world.

As *Undercover Economist* readers will know, Harford is an economist who applies a market-based way of thinking to all aspects of the world, and *Adapt* is no exception. He opens the book with a quotation from Hayek on knowledge ("The curious task of economics is to demonstrate to men how little they really know about what they imagine they can design") and comfortably asserts that "The market has solved the problem of generating material wealth" (12).

Harford sees the key characteristic of the market system as its evolutionary, decentralized, trial-and-error nature, with many firms pursuing their own goals in their own ways. And if individual firms fail, as they do in large numbers, that is not a problem: "The difference between market-based economies and centrally-planned disasters, such as Mao Zedong's Great Leap Forward, is not that markets avoid failure. It's that large-scale failures do not seem to have the same dire consequences for the market as they do for planned economies." (11)

The year 2011 is an odd time to make such a statement, even though he does follow it up with "The most obvious exception to this claim is also the most interesting: the financial crisis that began in 2007. We'll find out why it was such a catastrophic anomaly in chapter six". (11)

He goes on, "trial and error is a tremendously powerful process for solving problems in a complex world, while expert leadership is not. Markets harness this process of trial and error, but that does not mean that we should leave everything to the market. It does mean - in the face of seemingly intractable problems, such as civil war, climate change and financial instability - that we must find a way to use the secret of trial and error beyond the familiar context of the market." (20)

And Harford applies this remedy liberally throughout the remainder of the book: "The adaptive, experimental approach [trial and error] can work almost anywhere" (35) he says, and he applies it to climate change (chapter five), to business strategy (chapter seven), to individuals (chapter eight), and to financial collapses (chapter six).

So here is a summary of *Adapt's* remedy for uncertainty and complexity: "The three essential steps are: to try new things, in the expectation that some will fail; to make failure survivable, because it will be common; and to make sure that you know when you've failed." (36)

1.2 The contradiction

Theorem: *In systems with multiple levels, trial-and-error cannot be an optimal strategy at each level.*

Proof (by contradiction, and a bit like the Unexpected Hanging Paradox).

Consider an economy consisting of a government, a set of firms, and a set of employees. There are two sets of decision-making agents in this economy: the government must decide on a strategy for the economy as a whole and firms must decide how they will operate internally.

To keep things really simple, limit the available strategy space for each decision-maker to only two strategies: *trial-and-error* delegates decisions to the level below (government to firms, firms to employees) while *eggs-in-one-basket* mandates to the lower level what they must do. Then we can show that a recommendation to adopt trial-and-error both at the level of the economy as a whole and at the level of the firm, as Harford does, is contradictory.

If trial-and-error is the optimal strategy at the level of the economy, the government will adopt it, allowing each firm to identify and pursue a strategy of its own.

Some firms will choose trial-and-error, while others adopt eggs-in-one-basket. Which firms will succeed? There are two cases.

Case A: some trial-and-error firms succeed, and some eggs-in-one-basket firms succeed. In this case, trial-and-error is not the optimal choice at the level of the firm, it is simply one strategy among others (the others in this case being eggs-in-one-basket) that may be worth pursuing.

Case B: only trial-and-error firms succeed. In this case, trial-and-error is the optimal strategy for firms. But (here's the catch) if trial-and-error is the best strategy at the level of the firm, there is no need for experimentation at the level of government: the optimal strategy for the government is eggs-in-one-basket, mandating that all firms must operate in a specific (trial-and-error) fashion.

In short, if trial-and-error is the best strategy at the level of government, then we cannot say it is the best strategy at the level of the individual firm. And if it is the best strategy at the level of the firm, it cannot be the best strategy at the level of the economy.

1.3 Consequences

The contradiction appears in stories throughout the book itself and so undermines the author's argument. Here are two examples.

1.3.1 Capecchi and Tharp

Adapt tells us the stories of Mario Capecchi (pp 97–100) and of Twyla Tharp (pp 247–256).

Tharp is a brilliant and determined choreographer whose ambitious production *Movin' Out* was panned when it premiered in Chicago. According to *Adapt*, Tharp is a believer in the process of experimentation, filming hours of improvised dance in the search for just a few interesting moves. She treated the Chicago experience as a failed experiment and “made peace with her losses and immediately set about the hard work of winning back both the critics and the audiences.” (254) She reworked the show, and the result was a piece of groundbreaking dance that won her rave reviews when it moved to Broadway. An inspiring story.

Capecchi is a brilliant and determined biologist whose ambitious proposal to “make a specific, targeted change to a gene in a mouse’s DNA” (99) was panned when he submitted it to the NIH funding agency. According to *Adapt*, Capecchi is a “stubborn genius” who survived a remarkable childhood as a street urchin. He refused to treat the NIH proposal as a failed experiment, refused to make peace with his losses, and instead treated the NIH experience as an obstacle to be worked around, got grants for two other projects and used the money to carry out his panned mouse-gene project anyway. The result was a piece of groundbreaking research that won him the 2007 Nobel Prize in medicine. An inspiring story.

Harford uses the story of Twyla Tharp to promote the virtues of trial-and-error at the level of individual, and the story of Mario Capecchi to promote the virtues of trial-and-error at the level of the organization. But trial-and-error at the organization level means accommodating people like Capecchi, who is clearly not a trial-and-error individual and who exhibits precisely those traits (rejecting critics views, refusal to change) that Harford warns us against in the Tharp story.

1.3.2 Whole Foods, Timpson, Chile, and Wal-Mart

Adapt tells us the story of American high-end grocery chain Whole Foods and UK bric-a-brac retailer Timpson (pp 224–230). These two companies promote decentralization and experimentation, giving teams (in the case of Whole Foods) or stores (in the case of Timpson) independence from the company so that they can find their own successful strategies. Both exemplify Hayek’s ‘now familiar words “knowledge of the particular circumstance of time and place”’ (227). Harford uses their stories to show that “the world is

increasingly rewarding those who can quickly adapt to local circumstances” and to promote strategies such as delegation of power to the front lines of the organization. Peer monitoring “offers a subtlety and sensitivity that monitoring from corporate HQ simply cannot match” (229).

He contrasts the decentralized efforts of Whole Foods and Timpson with the failure of centralized planning, exemplified by “one of the most surreal examples of the planner’s dream” (69), Salvador Allende’s Project CyberSyn which looked to use a supercomputer to collect centralized reports of economic activity throughout Chile in the early 1970’s and to tune the economy. The project was “not a success”, and shows us “the way in which our critical faculties switch off when faced with the latest technology.” (70) The dream of “information delivered in detail, real-time, to a command centre from which computer-aided decisions could be sent back to the front line” persisted in the form of Donald Rumsfeld and his failed conduct of the Iraq invasion (chapter 2). The lesson is that “such [centralized] systems always deliver less than they promise, because they remain incapable of capturing the tacit knowledge that really matters.” (71)

Wal-Mart is mentioned only briefly in *Adapt*, on p 226. “Of course, this kind of business model [Whole Foods] is not the only way to succeed in the supermarket trade. Far more centralised supermarkets such as Wal-Mart in the US and Tesco in the UK are clearly very profitable”.

Wal-Mart owes its success to massive centralization of a kind that makes Project CyberSyn look unambitious. Famously, the Wall Street Journal reported in 2007 that “Wal-Mart’s centralized thermostat system in Bentonville, Arkansas, its corporate headquarters, actually uses a monitoring team to control the temperature for every store from this centralized location.” Centralization and scale permit remorseless cost-cutting precisely by minimizing local experimentation even in such minute decisions as store temperature.

Again, the contradiction plays out. Harford wants to argue that at the level of the economy as a whole and at the level of the individual firm, trial-and-error triumphs over misguided centralization. Yet trial-and-error at the level of the economy permits individual firms like Wal-Mart to pursue centralized, eggs-in-one-basket strategies, some of which turn out to be better than experimentation. You can’t eat your cake and have it too.

1.4 Centralized experimentation

Harford does note that Wal-Mart and Tesco “still experiment but have managed to centralise and automate that experimentation” (226), but gives no

further details. This sentence shows another failure of the book: a blurring of the line between experimentation (trial-and-error) and decentralization. Throughout most of the book he uses experimentation as a synonym for decentralization (tacit knowledge and all that) and is in favour of both, but sometimes - as here - he separates the two to make his argument fit.

The most dramatic case where he separates experimentation with decentralization is in the chapter on development aid (chapter 4). Part of the issue in pursuing a trial-and-error strategy is to identify successes – to design a “feedback loop” that permits successful ideas to evolve further – and chapter 4 looks at the use of randomized trials in development projects to provide that feedback.

Harford argues throughout chapter 4 that development aid projects can easily go wrong despite, or because of, the best intentions of those involved. To identify successful strategies he argues for rigorous experimentation based on clinical trial methodologies, and randomized trials. Why, I wondered throughout this chapter, does the use of randomized trials come up particularly in the area of development aid? I believe (and here I may be attributing ideas to Harford that are not his) that it comes from the usual economist’s idea that in development aid we must avoid woolly, sentimental thinking and adopt a hard-headed approach if we are really to Do the Right Thing.

The problem is, randomized trials are often not incentive compatible for the participants. In one example, Harford describes an experiment in Kenya to deliver a new set of textbooks to schools. The charity funding the programme “chose twenty-five schools at random” and distributed the books. The experimenters found, surprisingly, “little evidence that textbooks were helpful”. Chalk one up for randomized trials.

(I’m going a bit out on a limb for the next few paragraphs: if others with more knowledge of the subject can correct me, go for it).

What are the incentives at work at the level of the individual school? Given a choice between a set of new textbooks and no set of new textbooks, most schools would choose the books because the expectation was that they would improve the education experience. To succeed, randomized trials demand either a setup where there is no expectation of likely outcomes (as in another example he gives, an 18th century sea doctor treating scurvy with “oranges and lemons” or “cider, acid, or bring”), or where the subjects of the experiment are deprived of the right to choose. If there is a likelihood, going in to the experiment, that one option or the other is more likely to be beneficial, then it is in the interests of the experimental subjects to go with that option rather than participate in a randomized trial.

It is not surprising, then, that Harford can find stories of randomized

trials in the case of medicine and of development aid. Both scenarios rely on a powerless, voiceless set of experimental subjects. In scenarios where the subjects have a voice, randomized trials tend not to be tried despite their system-wide benefits because the incentives don't line up. Are there cases of trials in North American or British school systems where a random selection of schools get access to a new, potentially beneficial teaching aid? If there are, Harford has apparently not found them.

So here again the contradiction is at work. Trial-and-error at the level of the development project demands centralized control. The school textbook project demands that trial-and-error not be an option at the level of the individual school, for example. In this way, the message of the story is counter to the message he seeks to deliver.

Such contradictions occur throughout the book. I don't know much about development aid or randomized testing, but I do know a bit about technology. Harford gives Google a positive write-up for its famous "20% time" program for employees (pp 231-234), showing that it succeed by casting aside those projects that fail. But in demonstrating why failures need to be rigorously thrown aside (a "tight feedback loop") he writes "According to the TechRepublic website, two of the five worst technology products of 2009 came from Google - and they were major Google products at that, Google Wave and the Android 1.0 operating system for mobile phones. Yet most internet users know and rely on Google's search, Google Maps and Image search, while many others swear by Gmail, Google Reader, and Blogger." (234). By his own logic, Google should have abandoned the duds, and it did throw Wave overboard, but of course in the case of Android it persisted and now Android is the most widely-used smartphone operating system in the world. And Google Maps and Blogger, at least, are not products of the innovation program but were developed outside Google and then purchased.

Adapt also uses business professor Clayton Christensen's *Innovator's Dilemma* to bolster its case, in which Christensen shows how surprisingly primitive but cheap and "good-enough" technologies tend to displace advanced, high-end technologies. For the record, I found that to be an excellent book. But I first heard Christensen talk on the subject at a BlackBerry conference where he explained that the threat to BlackBerry was that its end-to-end design was vulnerable to a horizontal model that would not give the same highly-tuned experience, but which would do a good-enough job at lower prices. It sounded convincing, but within a year (I think) the real competitor turned out to be Apple's iPhone - an even more high-end, even more end-to-end design, and quite the opposite of the prediction. Apple, of course, is a hugely centralized company that puts all its eggs in one basket

when it releases new technologies.

The paradox of giving a recipe for success in any competitive environment goes deeper than Harford (or Gardner, see previous post on this blog) acknowledge. Any recipe, including trial-and-error, becomes self defeating. The key feature that evolution is context dependent, which both authors believe, undermines the remainder of this book.

1.5 The Financial Crisis

By this time, you can see that I think *Adapt* suffers from the very cognitive dissonance that is author warns against in his final chapter. In the face of contrary evidence, the author finds ways to accommodate the facts within his framework by stretching the argument in ways that are ultimately unconvincing. It's exactly this kind of flaw that Tetlock identified among his worst-performing experts (the hedgehogs). Nowhere is it more obvious than in *Adapt's* chapter on the financial crisis.

To identify successful strategies, Harford argues that “we should not try to design a better world. We should make better feedback loops” (140) so that failures can be identified and successes capitalized on. Harford just asserts that “a market provides a short, strong feedback loop” (141), because “If one cafe is ordering a better combination of service, range of food, prices, decor, coffee blend, and so on, then more customers will congregate there than at the cafe next door”, but everyday small-scale examples like this have little to do with markets for credit default swaps or with any other large-scale operation.

The chapter on the financial crisis misses the boat completely.

One part of the chapter deals with arranging incentives to encourage whistleblowers. There is only one problem with this, which is a lack of whistleblowers was not the problem with the mortgage market. There were people who saw the market coming, said so, and put money behind their words, and their stories have been chronicled in books like *The Big Short*. The problem was that such people were seen as Chicken Littles. The problem was that the “short, strong feedback loop” of the market was neither short nor strong: it sent fundamentally misleading messages, because that's what a bubble is.

The second idea *Adapt* has about financial crises is to provide greater visibility for regulators into problems and stresses in a system, and it is illustrated by a story about the final weekend in September 2008, as Lehmann Brothers collapsed. But again, that weekend was simply the final bursting of the pimple and, while replaying that weekend may have led to different

outcomes, the time was long past when the crisis could have been averted.

The final idea is borrowed from engineering systems, and has been floated again by others recently (Justin Fox here), and is to decouple various parts of the financial system so that failures in one area cannot spill over to failures in other areas. This may or may not work - I'm no expert - but I can say this: the proposal has nothing to do with the thesis of the book. The coupling between different parts of the financial system came about precisely because of an unwarranted belief in the virtues of the innovative, market-driven processes that Harford is promoting, and were governed by exactly the kind of feedback that he relies on (price, profits) elsewhere in the book.

1.6 Foxes and Hedgehogs

Like *Future Babble* (see previous post), Harford includes the work of Philip Tetlock early in the book, to show us not to trust experts. And that's fine. But he fails to see the depth of the paradox of expert failure: it applies to those who would replace experts as much as it applies to experts themselves.

Tetlock divided his experts into foxes (good at many things) and hedgehogs (good at one thing) and argued that hedgehogs are over-confident because they "reduce the problem to some core theoretical scheme' . . . and they used that theme over and over, like a template, to stamp out predictions". And that's exactly what Harford does here. He sees evolution as a fox-like strategy (trying many things and selecting a few) but doesn't notice that at the level of individual species, evolution gives us both foxes and hedgehogs, and both do perfectly fine.

Once the contradiction at the heart of the book is clear, it is not surprising that the book itself cherry picks examples where trial-and-error has succeeded, or where eggs-in-one-basket has failed. But such stories, while entertaining, make a notoriously shaky foundation for any kind of general structure, and so it proves here.

1.7 If not trial-and-error, then what?

So in the end, Harford fails in his attempt to sell trial-and-error as a panacea. It's easy to knock, I can hear you say, but do you have anything better to offer? Well probably not, but let me try.

While both Gardner and (especially) Harford place great emphasis on the usefulness of particular knowledge, and the need to recognize our limitations when it comes to seeing the future and to planning. "Allow room to experiment, to revise, and to adapt" is not bad advice so far as it goes. But

Harford pushes this argument too far, and so tumbles into contradiction. He seeks to use this lesson as a general, all-purpose lesson, which means that he is again failing to acknowledge our limitations when it comes to seeing the future and to planning.

We need to accept that there is no algorithm for success. In fact, any such recipe would be self-defeating. The process of achieving success is irreducibly specific, irreducibly individual, and irreducibly paradoxical. It is not the realm of science, logic and analysis - it is the realm of art, precisely because art is comfortable with paradox and self-contradiction in a way that science and logic is not.

Or: "If I knew the jazz of the future, I'd play it" as someone said.