Free Soloing the Future: How creativity could save the world (and why it probably won't)

Dan Holloway, Rogue Interrobang (https://rogueinterrobang.com)

Transcript of talk (as delivered) at Taylor Institution Main Hall, 23 October 2018 to celebrate the launch of Mycelium

Note: slides 1 and 2 not included in the accompanying downloadable Powerpoint for copyright reasons. Used on the night under educational exemption.

INTRO SLIDE

At 5.32 on the morning of June 3rd last year, Alex Honnold did what he does most days. He began climbing. Three hours and fifty six minutes later, our perception of what is humanly possible had been changed forever.

SLIDE 1

El Capitan in Yosemite Valley, California, is a wall of a kilometre of vertical granite, much of it so smooth as to be almost featureless. It is the centre of the climbing universe. The first ascent was made only 60 years ago, using countless bolts and thousands of feet of rope and ladders, and taking months.

Alex Honnold's ascent was made free solo.

SLIDE 2

That is to say, he made the climb on his own without ropes or any other form of safety or support. It was something everyone knew was impossible.

Like Honnold we also stand at a precipice. Unlike him, it is not one forged by nature over millions of years, but a precipice of our own making formed in just a few hundred. And while Honnold stood at the bottom, searching for seemingly impossible ways to climb, we stand at the top, scouring the face for seemingly impossible ways to prevent our fall.

Whether it's climate change, food security, where artificial intelligence will take us, the potential for nuclear annihilation, or the implications of genetic modification and transhumanism, we face a future that could so easily go catastrophically wrong. And yet these issues are so complicated we have very little idea where to even start in tackling them. It's this factor that lends them the name "wicked problems".

The one thing we do know with some degree of confidence is that our current way of doing things probably isn't going to help us find good solutions – after all, those ways of thinking and acting are what got us into this pickle to start with.

We know, in other words, that to help us solve our wicked problems we need creativity. And here, as throughout this talk, I am using the word creativity simply to describe new ways of doing things, new ways of looking at things, most of all new ways of putting things together.

And really, that should be the end of it. Or the start of it at any rate. We know we need creativity, so surely we have, as a species, embarked on a colossal quest to find and nurture every new idea? Well, no. And that's why the subtitle of this talk is so pessimistic.

But there's a problem. It's hardly a revelation – you will hear it in many different contexts, rattled off as an evident truth. But very rarely will people actually stop and think about what it means. Specifically about what it means every one of us needs to do if we are going to get around it and pull back from the cliff.

There is one simple trait that almost every institution, every profession, every company, every political party shares. Many organizations know they have this trait. Most of them say they want not to have it. And most of those say they have no idea how. The rest say simply that they can't. Or they pretend that they already have, while knowing that not to be true.

It is a trait that manifests itself in many ways. Every one of them is disastrous. Each of them is decidedly avoidable yet each can feel unavoidable. There are some very good reasons why groups have this trait, many of which come from the fact that for individuals it can be a really useful trait in most situations, and essential in many.

If human history were a classical tragedy, and heaven help us it can feel like that often enough right now, this trait would be our fatal flaw. Any observer watching this slow tragedy play out would see it at once, would watch as we agonised over this weakness again and again only to fail once more in the same avoidable yet inevitable way. They would watch as we walk knowingly into an endless night of our own creation.

The trait is this.

Progress is made by those who do something first. Fundamental (as opposed to incremental) progress cannot be made without something happening that has never happened before. Yet pretty much every structure, every system we have created is precision-engineered to excel at doing things for the second time. At an organisational level, creation happens by chance; it is only copying that happens by purpose.

Think of the mechanisms through which an organisation chooses which courses of action to support, and to what extent, as an algorithm designed to implement certain key goals given certain starting parameters. The problem isn't with the algorithm. That's dictated by the goals and the starting parameters. Neither is there a necessary problem with the goals that organisations have chosen. They are, after all, able to choose better ones. The problem is with the starting parameters. The rules about the environment in which organisations operate. Specifically the rules about which parts of that environment are fixed.

SLIDE 3

One of my other worlds, publishing, is a great example of this. Publishers all say they want to publish really exciting, fresh books. That's a great goal. But they also believe

that for a book to stand a chance of selling, it has to fall into a niche that's been proven to sell. The result?

SLIDE 3 ANIMATION

Lots of books about sparkly vampires.

At its most simple the problem we have is this. Creativity sees an inaccurate parameter as an obstacle and says "that obstacle doesn't have to be there. We can move it like this." And institutions keep coming back with the reply "No you can't. Because it does have to be there."

In the first part of this talk I want to talk about the skills involved in creativity, and what we can each as individuals do to train ourselves in them.

And then much of the second half will be about how we misidentify many parameters as fixed that are, in fact, flexible, with a view to discovering how we might do better, and as a result how our organisations might do considerably better at providing environments in which we can all use our creative skills so that solutions to the world's wicked problems might emerge, be nurtured, and flourish.

So what is creativity?

The two raw materials are incredibly simple. You need to know lots of things about lots of things. And you need to be able to use that knowledge to its full potential, which means being able to connect any one part of it up to any other part of it.

There are some fascinating studies from neuroscience that help us pin down these ingredients and offer suggestions as to how we might develop them. The world's leading expert on the neuroscience of creativity, Nancy Andreasen, has shown that people who exhibit more of the traits we associate with creativity have higher levels of activity in their association cortex areas. A more recent study by a group of scientists in the Netherlands, which happened to contain leading memory athlete Boris Konrad, looked at the effect on the brain of training using the techniques employed by memory athletes.

SLIDE 4

They specifically looked at the kind of visual techniques we might call the memory journey or mind palace. They found this training heightened activity in those parts of the brain Andreasen had LINKed with creativity. So memory, and one particular approach to memory that we will explore later, is clearly important.

But it is not enough. Another famous study carried out by Charles Limb looked at what happened to rappers when they began to improvise – that is, when they started spontaneously generating ideas the way we think of people engaged in creativity generating ideas. Limb's MRI study found that the moment the rappers started freestyling, their frontal lobes, the part of the brain linked to self-censorship, shut down. Similar results were found in those other great improvisers, jazz musicians.

So being creative is a combination of learning things and being able to form connections between them (or maybe, we should say being able to turn off the mechanism that stops us exploring connections between them).

In many ways that sounds simple, and rather obvious. But if that's the case, a lot of our thinking about creativity fails to reflect that obviousness. Especially when it comes to what it means to "know lots of things."

What I mean is this. We place too much value, and we always have done, on memory. There are obviously good reasons for this which go back to the earliest days of life, and forward from there to the campfire when you would be really hacked off if Homer got to the bit where Hector and Achilles square up to each other only to say "sorry, folks, I've forgotten what happens." But from educators to executives we still, as a society, place inordinate emphasis on remembering what happened next and what goes where.

It's somewhat reminiscent of the kind of muddled thinking that has people on social media drooling over a picture of a group of folk with their heads in a book

SLIDE 5

while a group of people with their heads in an even tinier object that gives them access to every book ever written has people shaking their heads in despair.

SLIDE 5 ANIMATION

You don't need to remember how long a manatee's fins are! What you do need to know is just how many things you can do with that information once you've looked it up on Wikipedia. Or Sci Hub.

All of this isn't to say that memory isn't essential to creativity, because it is, and in a moment we'll have a look at some ways to make your memory better. It's rather that placing too much importance on memory for memory's sake is a symptom that we don't quite realise just how powerful knowledge can be.

Let me explain.

Thinking about knowledge just in terms of what we can recall means that we think of it something like this.

SLIDE 6

In this diagram, our knowledge is the sum of the things we know. Every time we learn something, we add something to that sum. But the more we know, the less, as a proportion of that total, each new thing adds to it.

But imagine instead we thought about our knowledge not as the sum of the things we know, but as the product

SLIDE 7

as the connections between bits of data and not the data themselves. Not only does this seem to fit nicely with Andreasen's results. It paints a picture of knowledge as a thing of immense, approaching infinite potential. And something that, far from offering diminishing returns with each new piece of data added, increases in power by orders of magnitude each time we learn something new.

Rethinking the relationship between knowledge, memory, information, and creativity in this way turns the way we think about human potential on its head. Showing us the power that creativity could unlock offers us hope for the future and an incentive to pursue that future. It fundamentally shifts our perception about creativity and that, as we will see later, is the essential first step in starting a creative revolution.

There is a further problem that comes with our overemphasis on memory as the foundation of our knowledge. Suppose we accept that the real value of knowing things lies in the connections we can make between them. If we have grown up believing in the priority of memory, the chances are that the way we learn things will stop us connecting them.

Let me show you what I mean. I first became interested in memory, specifically in improving it, while studying for my finals here back in, well, let's just say it was last century. Like many people my age, my note taking was revolutionised by reading Tony Buzan's Mind Map book. Mind maps are great ways to learn things.

SLIDE 8

Especially the kind of things you find on Theology and Philosophy course in Oxford. When it comes to remembering which school of early Christian thought held the logos-sarx view of the incarnation of Christ and which the logos-anthropos model they are fabulous.

There is a lot I could say about what makes mind maps actually less than perfect for learning (for one, Buzan's "they look like the structure of distributary systems in the natural world" sounds not unlike the memes you see that claim that walnuts are good for the brain because they look like the brain, and so on). But the real problem is the one they share with almost every other way of learning. They help us to learn things by putting those things in their "correct place". They fix things.

SLIDE 8 ANIMATION

Anchor them. Isolate them at the end of a genus-species line like an obscure category of metaphysical fiction in an Amazon shopping recommendation. Every trick these techniques employ to help you remember a thing simultaneously hinders you from imagining that thing from existing anywhere but its "proper place".

You simultaneously learn the thing and learn not to be able to use it.

Instead of learning in this way, we need a way of thinking about the things we know that not only doesn't make it harder to use them creatively but helps us to do so. This might also have benefits for learning things for exams, where what matters is not knowing something (as with creativity, what you know is simply the data set you get

to manipulate) but understanding it – and understanding is essentially the same kind of iterative process as creativity, constantly shifting in perspective between the macro and the micro scale, asking alternately where a new piece of knowledge sits within a much wider conceptual framework and how your wider frameworks must be tweaked in the light of new knowledge. This continual zooming in and out, which is rather like television's first psychedelic special effects, is an incredibly effective way both to learn, and to learn from what you learn.

In other words it really matters HOW WE LEARN THINGS

A much better way of learning things than mind maps is the mind palace.

SLIDE 9

You have probably come across the idea through Sherlock, or Hannibal. It's the technique the memory athletes in the study we looked at earlier use, and it goes back thousands of years, to Cicero and beyond. In short, the technique is this. To remember a new piece of information, we link it to something we already know. In fact to something we know so well that it takes no effort to recall it at all, such as the route we take to work each day, or our home, or maybe the Bodleian Library. What we do for each piece of information is to find a sensual way of imagining it, usually a picture, which we make as memorable as we possibly can by exaggerating it, bringing other senses into play, making it humorous, or rude. And then we attach it to, say, the statue of William Herbert.

SLIDE 10

So say I wanted to remember that the Council of Chalcedon was held in the year 451. Here's William Herbert, appropriately outside the Bodleian's Divinity School. I add a vivid image of the Council of Chalcedon taking place.

SLIDE 10 ANIMATION 1

And to remember the year 451 I make the image burst into flames, reminding me of the book Farenheit 451.

SLIDE 10 ANIMATION 2

Now, when we want to recall that piece of information, all we have to do is travel in our mind to Herbert's statue and discover what it was we left there.

We know this is a powerful way to grow the part of our brain associated with our usable knowledge. Mind palaces employ imagery, and they employ the principle of association, and both of those are incredibly important.

The reason we don't all use them already the way people in the ancient world did has, on a mundane level, to do with places like the Bodleian. The more we could turn to books to remind us of things, the less we needed to be able to remember them. But it also has to do with the rather extraordinary arguments that embroiled 16th century Europe. The use of the imagination – of images – was considered deeply

suspect to the extent that using these techniques could see you accused of summoning demons and even getting burned at the stake.

SLIDE 11

It was much safer to learn things using lists formed from words, lists and words that mapped the place each thing had been assigned in the natural order of things.

And so we find ourselves having to learn how to create mind palaces afresh.

But although mind palaces are great, they still have some of the problems of mind maps and lists,

SLIDE 12

so they are not a magic bullet. By their nature, they isolate things, they're about the rooms and locations where things are placed.

SLIDE 12 ANIMATION 1

If we are going to be free to use our knowledge effectively, what we need, in essence, is a mind palace that has leaky plumbing! Once you've "learned" that something "belongs here" it's very hard to shift in your head to "it might belong there", so we need ways to think about things that mean they could go here but they could just as easily go there.

There are very good reasons why we use shortcuts to help us categorise things. Like the self-censoring mechanisms of our frontal lobes such shortcuts have been useful for us as a species. They have freed up our time so that we have been able to contemplate other things.

Of course, those subsequent thoughts and everything that follows from them has led us beyond being able to meet our needs in times of famine as well as feast to a place where we are in danger of destroying ourselves. Only now we have finding shortcuts baked into the way we work. Just at the time where we need to be thinking of everything in longhand.

When it comes to how we think about things, instead of a minimum kit that helps us place things, what we need is a maximum kit that makes it as easy as possible to move them around. When we think of things we need to get used to thinking of all of their properties. All of their associations, personal, cultural, historical as well as physical. These properties are hooks, waving around waiting to latch onto one another. The more of these hooks you can set, the more connections you will be able to find with other objects.

Take the kind of question you might get asked in a standard creative thinking test.

"What would you get if you crossed a dog with a skyscraper?"

SLIDE 13

Our shortcut for identifying a dog might include for example being a four-legged furry mammal and so on

SLIDE 13 ANIMATION 1

But when it comes to answering the question, that isn't very helpful. If we maximise our dog definition kit though so it also includes that dogs are known for their friendship of humans, for giving their Latin name variously to puncturing teeth (canines) and a group of islands (The Canaries),

SLIDE 13 ANIMATION 2

Then we start getting some interesting possibilities

SLIDE 13 ANIMATION 3

Obviously, this is something we need to do sparingly – we need shortcuts to be able to do anything at all beyond sitting all day contemplating. But it is nonetheless an essential habit to cultivate, because the risk of spending too little time being distracted may be greater than that of spending too much.

I developed Mycelium, which is the reason why I'm actually here tonight, as a way of pulling all these things together into a game format that would take a few seconds to learn. The object of the game is simple.

SLIDE 14

You use two decks of cards to set you a creative thinking problem. One deck of cards has questions about how to link two unconnected items together. The other has 40 different items and you choose two to fill in the blanks. Everyone playing then has 5 minutes to come up with as many answers as possible. The more people who come up with any particular idea, the fewer points that idea scores. This format works on expanding your knowledge across a whole range of unconnected subjects, and by rewarding the most outrageous answers, you are effectively training your frontal lobe to turn off, giving the brain a small dopamine hit every time it refuses to self-censor.

I want to spend some time talking about the barriers to creativity. I want to explore why we, and by we I mean humanity, society, business, academia, and probably most of us in this room, are so bad at being creative. By that I don't mean, as speakers at TED talks inevitably mean when they say with a whoosh of inspiration "but we're all creative!" that we as individuals aren't creative. I mean something far more important – that we, in our groups, are very bad at acting on creative ideas.

Understanding what we do badly is more helpful than starting by looking at what we do well for two reasons. The first is deep practice. This has become somewhat of a buzz phrase, which is unfortunate, because it's an incredibly valuable concept. It

came out of studies by Anders Erikson of hot spots of outstanding talent in particular activities, from tennis to the violin to skateboarding.

SLIDE 15

What Erikson found was that people in these hotspots all practised the same way. Instead of running through the same exercises over and over, they only ever focused on practising the things they couldn't quite do. Once they had got the hang of some part of the activity they moved on to something else. By only focusing on the things they couldn't do, they were always at the edge of their capabilities, and this turned out to supercharge their progress.

The second is precisely the kind of bias that we will encounter here. Most people, when asked what they have done well in an area, will alight on things that have little to do with skill and much to do with luck, or will confuse necessary and sufficient cause (to succeed in a task you will have to put in a lot of work, but putting in a lot of work is no guarantee of success). In part this happens because we like to believe we are responsible for our success – we like to feel we deserve a pat on the back (as opposed to the deep practice approach that seeks out the things we have done badly or the areas in which we have made bad choices so that we can improve next time). But it's more than our egos getting in the way. It's also that we lack the criteria to make good judgments about what we have done well until we have scrutinised the biases that lie behind those criteria. That's another way of saying "we know that however we make decisions, that process has landed us with climate change, unregulated artificial intelligence, and vulnerability to pandemic threat, so maybe it's best not to use the same process to decide what we did really well." Understanding some of the flaws in our decision making processes might not lead directly to being able to save the world. But it might help us mark some of our potential paths as unsuitable for use.

So, if creativity has the potential to pull us back from the cliff's edge, why am I so pessimistic?

The first reason is expressed in a phrase that's often repeated in the business world. **What got us here won't get us there**. What that means is simple to the point of being banal. The habits, assumptions, and activities that have landed us with each of the wicked problems we face are not going to help us overcome them – left to our own devices we will inevitably keep on driving straight over the cliff edge.

The reason why this should be is fairly obvious. We didn't set out to destroy ourselves. That we are on the verge of doing so is because the assumptions that have driven our collective actions have had some fairly nasty unexpected consequences.

We know we need to try to address these problems. Again, very few people – openly at least – look at looming disaster and shout "bring it on". But the people in the best position to empower those with possible answers are, inevitably, those same people and institutions whose assumptions and conceptual frameworks caused the problems – despite their best intentions. So it would be the height of nonsensical magical thinking for us to presuppose that however much they genuinely want to

solve those problems their attempts to do so will achieve anything but an inevitable cliff dive.

It is hard to see a way out. Because these disastrous assumptions, whatever they might be, have become institutionalised, we no longer think of them as changeable. They have taken on the status of laws or, in classical Oxford speak “the way things have always been done.”

So those who adopt creative approaches to problems, ones that might work, will always have their ideas overlooked. It can’t be otherwise – they are breaking the rules of “how it’s done”! So they will be able to pursue them only despite and never because of, the institutions and individuals whose job it is to nurture and enable possible solutions.

Recruitment is a great example of this with its focus not on the skills actually needed for a task but on proxies, which end up becoming so fixedly the focus of HR teams that people who have the actual skills to do a job will never be given the job because they lack the relevant proxy. Seen from the outside it is somewhere between Kafka and a cartoon. Form the inside the logic is flawless.

We see it in academia in the system of peer review and grant awarding. And we see it anywhere a marketing team has any say over product development.

In every one of these ideas what is happening is the operation of the tragic human trait we saw at the very beginning. The solutions to problems that have been created by following a set of rules will most likely come from adopting different rules. But the first rule of “what got us here” is “replicate what has been shown to work.” Because the answer to a problem is new and has not been shown to work (if it had, the problem would be gone!), this rule will never allow us to select it form the set of possible solutions. We will benefit from it only if it comes from a tiny pool of the independently wealthy – rather than the vast pool of available people with brilliant ideas.

The second problem is an extension of the first. The term I’ll use for it is **the local maximum problem**. It relates to how we go about finding the best answer to any question and it can be explained very simply by showing a graph.

SLIDE 16

The problem is this. How do we go about finding the highest point on any landscape? Most means of solving the problem start with our current position, and scan the immediate landscape for somewhere higher. If you find a higher point, you move up, and you continue this process until there are no more higher points.

Looking at this in the context of a map shows just how bad it is as a technique, and yet when it comes to real world problems, it’s often the only way we can think of. We see it in product development with iterations of the same device, or training methods in sport – looking for the best model by starting with one that works reasonably well and employing techniques that find tiny improvements to it until it can’t be improved any more.

Again, this is clearly ridiculous. But trying to figure out how to do things differently so that we find the actual highest peak, as it were, is incredibly hard. If we had an infinite amount of resources we would run every possible scenario for an infinite length of time and see which one comes out on top.

But no one has infinite resources. So we have to find a way of deciding which scenarios to run and which to avoid – and yet the only reliable way of deciding is running all the scenarios. And so we come back to rules. In fact we come back to the same rule we always come back to – repeat what we know works well. And the only way we can improve on this at all is to find our own version of the technique run by the longest-running and widest experiment of all – evolution. That is to say "repeat what works but add some randomness". When I speak to people working in computer modelling, this is essentially what they do, and they spend most of their time thinking about the most effective ways of adding randomness.

When it comes to trying to solve wicked problems, we don't even do this much – there is no wholesale investment in random projects, no lottery for research grants, no sticking a pin in a pool of applicants for research posts or a list of applicants for panels at think tanks. Failure to introduce even this tiny element of randomness is essentially to bet everything on whatever slope we find ourself on being a nursery slope of the highest hill anywhere – even though we know it's also a hill that has a cliff we will all fall down if we keep fumbling around.

I want to end by looking at one last bias barrier. Because although in many ways it's the most frustrating of all, it's also one that offers some hope. Not hope that our institutions will suddenly discover a strategy to nurture genuine creativity, but hope that there might be a possibility of pushing them, step by step, to a place where they might.

The Overton window is a concept that comes from politics.

It's a very simple idea that explains much of the frustration we might feel with the political landscape. It states that for any political question, there are many possible answers, that come from many different points across a vast spectrum. But only a very few of those possible answers has any chance of being adopted, because people will only consider things that lie on a very small part of the possible spectrum, just like what people see as light covers a tiny fraction of the wavelength spectrum.

We might consider for any spectrum, we can overlay a line that looks something like this.

SLIDE 17

The only answers we will ever choose to a problem are those which lie somewhere in that window. At the very edge of the window are the kind of ideas you might find expressed in the kind of viral memes you find shared along with a suitable mix of emoji and exclamation points. If you place the middle of the window over the point on the line that represents what most people think, you will get an SENSE of which ideas lie inside the edges of the frame and which lie outside.

By now you can probably see where I'm heading. The behaviours and ideas at the centre of the window, the ones we take for granted, our norms, are the ones that have paved the way to the cliff. Creative answers lie so far along the line they are probably somewhere between fuck that and you're having a laugh.

SLIDE 17 ANIMATION 1 & 2

The future of work is a great example of this. It's kind of obvious that any non-dystopian society by the year 2050 will have automated just about every job and will pay a guaranteed living income to all its citizens regardless of what they do or don't do. But no politician today can even entertain putting such an idea to voters because it's so far outside the Overton window.

The Overton window is a really useful way of explaining why truly creative ideas are so hard to implement. Because to be creative is, essentially, to be Cassandra, the prophet from Greek mythology who was cursed always to tell the truth but never to be believed.

We have already seen that truly creative answers to wicked problems must come from outside the disastrous streams of thought that created the problems in the first place. They must, in the language we have been using, avoid the local maximum to find the true summit. If what got us into this mess won't get us out of it, creative pickle solving must come from another perspective.

All of which means that it is baked into creativity's DNA that the solutions it offers come from outside the Overton window of the problem being addressed. Anything inside the window can't be creative. Anything outside the window won't be accepted by the desired audience.

So why did I say this makes me optimistic when it sounds pretty much like Terminator crossed with the Fall of the Roman Empire? Well, because analysing the creative answer to wicked problems in this way shows us something most of us usually miss. Our creativity has not one task to perform. But two. That is not how it should be, of course, but it's how it is, and if we recognise that much then we can use the information and get to work.

What the Overton window shows us is that if we want to improve the world, we don't shout loudest about what we believe the answer to be. That will get us nowhere. In fact it may be counterproductive, because all the time, people saying the opposite, provided they keep just inside the window,

SLIDE 17 ANIMATION 3

will be pulling opinion further away

SLIDE 17 ANIMATION 4

from possible answers because the only way to balance out that sideways pressure is to exert a counterpressure at the other edge of the window, which we are failing to do.

SLIDE 17 ANIMATION 5

So finding potential solutions to wicked problems is only the first of our tasks. The second is to apply constant and gradual pressure to the edge of the frame so as to shift the window of what people will consider, ever so slowly, until eventually it has shifted enough that our solutions lie at its edge. What that means is this.

We might consider it the job of our institutions to nurture creativity. But they will not. Not from bloody mindedness or malice but because, for the reasons we have seen, they cannot. It is our job to perform a second task – that of bringing them to a place where they can.

This schematic of the terrain around the Overton window can be used to illustrate the two ways we might do this. The first is what we might call the 1000 cuts strategy. If we can find ways of constantly nudging the edge of the window, eventually we will pull it to the place where we need it to be.

SLIDE 17 ANIMATION 6/7

The alternative is what I call the big ripple. Some events are so big that the ripples they cause reach the people whose opinions are at the edge of the window, and they are then inspired to follow where new ground has been broken, creating movements, normalising what once seemed fantastical.

SLIDE 17 ANIMATION 8

Which of course brings us back to where we started. Alex Honnold’s ropeless ascent of El Capitan was one of the most remarkable big ripple achievements I can remember. Most meaningful shifts need both the big splashes and a relentless flow of small drips. Our long path away from the precipice is no exception.

By definition only creativity can rescue us from the future we have brought on ourselves. But it will only do so if we remember that finding solutions to our wicked problems is only the first step. Using our creativity to find a receptive audience willing to implement them is the hardest part of all. The good news is that it is simple, and fun, to become more creative. And that tiny steps are as important as giant leaps in the journey ahead of us.

SLIDE 18