



Achieving change that lasts

True and long-lasting change only happens if the structure of the underlying system is changed. **Costas Markides** and **Anita McGahan** explain why.



On 30 May 2008 a terrible hit-and-run incident was caught on CCTV in Hartford, Connecticut. It showed 78-year old Angel Arce Torres crossing a two-way street when he was hit by a speeding Honda. The amazing thing (captured on the video) was that nobody bothered to dash to his help or even stop the traffic as he lay motionless in the street. Several cars passed by him without stopping while a few people stared from the

sidewalk. Some approached Torres, but most stayed put until a police cruiser responding to an unrelated call arrived on the scene. The incident provoked an outcry in the press and the Hartford Chief of Police Daryl Roberts was quoted as saying that the lack of action on the part of passers-by showed that the nation had lost its moral compass. "At the end of the day we've got to look at ourselves and understand that our moral values have now changed," he said. "We have no regard for each other."

It is interesting that the police blamed our "changed" moral values for this state of affairs as if an incident such as this one could not have happened half a century ago because – presumably – we had stronger moral values then. Yet, a very similar incident did take place 50 years ago. It happened in New York City in 1964 and involved a 29-year old woman called Kitty Genovese.

Genovese had driven home from her job working as a bar manager late on the night of



13 March 1964. She parked her car about 30 metres from the door of her apartment, which was at the rear of the building. As she walked to her apartment, she was attacked by a man who stabbed her twice in the back. Her cries for help were heard by several (sleeping) neighbours and at least one of them shouted at the man to “leave that girl alone”. When the attacker ran away, Genovese was able slowly to make her way to the back entrance of her apartment building, where she then collapsed. The attacker returned ten minutes later and found her lying seriously injured in the hallway there, where a locked doorway had prevented her from entering the building. He proceeded to kill her.

The murder caused a sensation because a *New York Times* article published two weeks after the event claimed that several people had heard or observed portions of the attack, yet nobody did anything to intervene. According to the *Times* story, 38 people witnessed this

murder but did nothing. Police later revealed that the figure was most probably more like a dozen people and that none saw or was aware of the whole incident. Many claimed to have heard the cries but were not aware that a homicide was in progress. Although the initial perception (that many people had observed this murder and still chose to do nothing) is probably inaccurate, it remains the fact that few of the New Yorkers who did hear the cries actually called the police and nobody actually went downstairs to investigate or to offer their assistance.

Some people have argued that we may not stop to help because we are afraid that we may get in trouble by doing so

Why didn't people help?

What could possibly explain why rational human beings did not do what all of us consider obvious (i.e. helping fellow human beings in need)? As noted above, some people have pointed the finger at the erosion of moral values in our society. Others have claimed that our culture is to blame for encouraging individualistic behaviour and the pursuit of personal benefit with little concern for the common good. Yet others have argued that the problem lies in the time pressures that we are all under, which shift our priorities. Even the underlying incentives have come under criticism: some people have argued that we may not stop to help because we are afraid that we may get in trouble by doing so. After all, we have all heard of the case of the man who rushed into a burning house to save a woman, only to be sued by that woman for damaging her spine in trying to get her out of the house. ➤

The key point to notice here is that there are several factors that produced the suboptimal behaviour – incentives, time pressures, underlying moral values, group-think, diffusion of responsibility, fear of excessive responsibility and so on. Incentives are important but they are only one of these factors. This implies that fiddling with the incentives alone will not be enough to change how people behave. All the underlying factors – or at least a lot of them – would have to be influenced if we are to change the behaviours they produce. Failure to ‘correct’ the underlying factors will lead us to the same suboptimal behaviours again and again, no matter how much regulation we put in place or how much we change the people who produced the suboptimal behaviours in the first place.

Numerous celebrated social psychology experiments have determined that the collection of these factors (known as the ‘situation’ or ‘social context’) can have a bigger influence on how people behave than our own personality or moral values. According to one estimate, up to 70 per cent of the variance in our behaviour can be explained by the situation we find ourselves in (with our personality accounting for the remaining 30 per cent). The implication of this finding should be clear: if you want to change how people behave (such as getting them to stop and help people in need), you have to change the situation they find themselves in.

Academics in another field – system dynamics – have come to describe the collection of factors that influence how we behave in a different way. To them, all these factors make up the ‘underlying structure’ of the system in which we operate. The ‘system’ may be mechanical, electrical, economic or social. Whatever it is, system dynamics argues that “the underlying structure of a system creates the behaviours in that system.” As a result, actors within a system will behave in ways that are conditioned by the structure of that system. In a sense, we are all prisoners of the system.

This implies that if people ‘misbehave’ in a system, the last thing we want to do is to blame the individuals. Instead, we should change the underlying structure of the system that produced the suboptimal behaviours in the first place. All the factors that we have been talking about so far – such as incentives, moral values, time lags, time pressures, groupthink and so on – are all part of the underlying ‘structure’ of a system.

Beer lessons

Nowhere is this principle more obvious to see than in the story of the ‘beer game’, an educational exercise (originally known as the

‘production-distribution game’) developed by Professors Jay Forrester and John Sterman at MIT. The game is played on a board that represents the production and distribution of beer. Each board consists of a team of four players: Retailer, Wholesaler, Distributor, Factory. In each simulated week, customers purchase from the retailer who moves the beer (represented by pennies) out of their inventory and then places an order with the wholesaler. The wholesaler ships the beer requested by the retailer and places an order with the distributor. Likewise, the distributor ships the order placed by the wholesaler and then places an order with the factory. The factory ships the order placed by the distributor and then manufactures beer from its raw materials. At each of these stages, there are time lags, representing ordering and shipping delays.

The game is supposed to be played for 50 weeks and the players’ objective is to minimise total team cost. There are two types of cost in the game: inventory costs 50¢ per beer per week, while an order backlog (i.e. inability to deliver) costs \$1 per beer. Thus, the players have to balance two costs: they need to order enough beer so that they can supply from inventory whatever is ordered from them but not so much that they have expensive unused inventory. The only other rule is that no communication is allowed between players.

We have now played this game with hundreds of executives from different countries, ranging in age from 25 to 50 and in seniority from young managers to board directors. No matter who plays the game, the patterns of behaviour are the same each time. The most prominent behaviour to emerge every time is the suboptimal way of playing the game even by seasoned managers. For example, people place such unreasonable orders for beer and overreact so often in doing so that team cost at the end of the game can be anything between \$2,000 and \$10,000. Very seldom does a team achieve a cost of less than \$1,000. To put these results in perspective, the optimal performance for the game is well below \$500. So, actual performance is ten to 20 times worse than optimal!

There are other commonalities in the behaviour of people. For example, there are always fluctuations in the orders placed as well as the inventory held – both go up and down like waves. And the fluctuations amplify: the ups and downs in the orders placed increase steadily down the supply chain from customer to factory. For example, if the retailers increase their order by two beers, by the time the shock is transmitted down the line, the factory is increasing their order by up to 20 beers!

To realise how incredible this behaviour is, all you need to know is that customer demand



(i.e. the demand faced by the retailer) is actually kept constant through most of the game. It starts at four beers a week, rising to eight in the fifth week and remaining constant thereafter. When we reveal this to the players, disbelief and despair follow. How is it possible, in a situation that is so (artificially) simple, that people always behave so suboptimally?

Once everybody accepts that their behaviour is indeed suboptimal, we try to understand why they behaved in such a way. The first reaction of people is to blame other team members for overreacting. Many others blame the game itself – it’s not clear or it’s lacking information, they say. However, the important question to concentrate on is why is it that no matter who plays the game, we always get the same suboptimal performance, which is characterised by fluctuations, amplifications and phase lags?

It soon becomes apparent that there are certain structural features of the game that induce every player into certain predictable (and suboptimal) behaviours. For example, the time delays in ordering and delivering beer induce people to overreact, or even panic, since it takes time for any inventory or backlog problems to be corrected. Similarly, the multiple layers in the system amplify disturbances (such as an increase in orders at the retailer end) and create shock waves down the system. The situation is made worse by the lack of communication among players: since they have no information about what other players are doing or the pressures they are under, they forget all about team costs and try to optimise their little piece of the system. Unfortunately, forgetting to account for the impact of decisions



Changing a system requires a much more sophisticated approach than only changing the individuals involved

and actions on team mates is disastrous because of all the interrelationships in the system.

The message that soon emerges from a post mortem is that it is the underlying structure of the game that creates the suboptimal performance of individuals: the time delays, the lack of information, the inability to forecast, the multiple layers in the system. The implication is that if these structural 'faults' are removed, behaviour will improve. Recent research has demonstrated that this is exactly what happens. Behaviour – and therefore results – improves considerably when intermediaries and ordering delays are removed and when all parties have access to customer demand data. As if by magic, 'inefficient' people are transformed into super-efficient managers!

Fully appreciating how the underlying structure of the beer game creates the behaviours we observe is one of the profound lessons of the game. Equally profound is the realisation that the underlying structure is much more than incentives. And the implication that becomes clear from all of this is that blaming or punishing people will not improve behaviours in such a system. Rather, behaviours will improve only after we 'fix' the structure of the game. And the changes required are nothing major or earth shattering – just a few small

steps on our part can have profound effects on how everybody behaves in the system.

Changing the people without changing the system will achieve nothing

Though simple, the principle we are proposing is actually hard for people to appreciate. The moment something goes wrong, we immediately search for scapegoats. We have this inbuilt belief that if someone did something wrong, this individual needs to be punished for it. We rarely look to rectify the underlying structure of the system that encouraged that individual to misbehave.

For example, while accepting the 2010 Oscar for best documentary, Charles Ferguson, the director of the movie *Inside Job*, lamented: "Three years after a horrific financial crisis caused by massive fraud, not a single financial executive has gone to jail." This is a classic example of blaming people for bad behaviour without paying enough attention to the underlying reasons that led them to the bad behaviour in the first place. The sad truth is that if the underlying reasons are not corrected, the same bad behaviours will emerge again and again no matter how many people you punish or how severe the punishment.

The same bias manifests itself in our attempts to improve behaviours in a system by removing the bad apples and replacing them with 'better' people. This almost always fails. It doesn't matter how "good" the new people are. If you put them in the same situation or structure, you will soon get the same bad behaviours out of them as you did with 'bad' people.

For example, in April 2003, the US Securities and Exchange Commission filed a complaint against Henry Blodget, an investment analyst with Merrill Lynch, alleging that Blodget committed securities fraud and violated the rules of the stock exchanges on which his firm traded. Blodget's website defended him by arguing that he was 'keelhaunched' by Eliot Spitzer, then the Attorney General of the State of New York, who later underwent his own keelhauling in a sex scandal and was replaced by his lieutenant, David Patterson, who was later accused of inappropriate behaviours of his own!

How did Blodget, Spitzer and Patterson defend themselves? None was very successful. Blodget settled and paid millions of dollars in fines. Spitzer left government permanently. David Patterson declined to seek election to the Governorship. Ironically, in each instance, 'culture' was identified as the culprit. Spitzer in particular identified the culture of Wall Street as at the heart of Blodget's behaviour. Similarly, when Spitzer got into trouble, the culture of political leadership – the same culture in which Bill Clinton had an affair with Monica

Lewinsky years earlier – was identified as the problem. Yet in each instance, the individuals charged paid huge personal penalties while the culture was left untouched.

We hold individuals responsible for systemic problems because privileged leaders – those in charge – represent the system itself. We assume that the best mechanism for changing a system is to lead from the helm and that failure to change the system by a leader is sufficient grounds for keelhauling. But changing a system requires a much more sophisticated approach than only changing the individuals involved. Replace the people without changing the values, culture and incentives and the same mistakes will occur over and over.

What will it take to prevent bankers and analysts such as Blodget from succumbing to the conflicts of interest in their professions? The challenge requires thinking differently about the problem – a task made all the more difficult because it requires creativity at precisely a moment in time when the mandate for change is urgent. Changing the culture of banking requires thinking differently about why Blodget and other analysts recommended stocks they didn't believe in. It requires understanding why bankers sometimes make loans that they know can't be repaid and why consumers often borrow money without thinking about the consequences.

Changing the incentives without changing the system will achieve nothing

The same idea applies to incentives. We have an inherent bias that to get people to behave in ways we want, all we have to do is to give them the right incentives. This rarely works because incentives are only one element that makes up the underlying structure of the system. Even when incentives produce short-term change, they very rarely produce the lasting change in behaviours that we want. Only a change in the underlying structure of the system will achieve this.

Consider, for example, our attempts to get pharmaceutical companies to sell more of their life-saving drugs in poor, developing countries. In 1994, the World Trade Organization created an incentive for pharmaceutical companies to sell drugs in developing countries. Some 17 years later, there is little evidence that the incentive has worked. What went wrong?

To answer this question, it's important to understand how the incentive worked. It was based on enforcing patents on the drugs in poor countries (more precisely, in every country in the World Trade Organization). The way that patents work is that they give the seller a monopoly on the drug for a number of years – usually about 20 – after its invention. Under >

the new policy, the sellers would be assured that imitators would not compete down profits even in poor countries, some of which had a history of hosting imitators.

The policy was adopted on the theory that the incentive (i.e. the patent) would make it worthwhile for drug companies to distribute their drugs in poor countries. The goal was to create benefits that would help poor people as well as the companies themselves. The most important benefit in developing countries would be that drugs would become available. If the incentive worked, then pharmaceutical companies would have an incentive to develop new drugs that were specifically designed for the kinds of diseases that turned up in developing countries – such as malaria, tuberculosis and HIV.

Nearly 20 years later, it seems that the incentive has not worked in the way that was envisioned at its creation. There was not a flood of drugs into poor countries with the implementation of patent protection. One issue relates to ability to pay: if the recipients of the drugs cannot afford the drugs, then there are little profits for the drug companies in distributing them internationally. And without the prospect of a return the drug companies also do not have much incentive to develop drugs for the diseases of the poor.

A closer look at the situation might cause you to raise your eyebrows. Sometimes,

agencies and organisations, such as the Gates Foundation and even the US government, pay for the drugs. This means that drug companies get paid even if the patients themselves cannot afford the drugs they need. And some drug companies sometimes offered to charge low prices out of humanitarianism. Generic manufacturers made available versions of drugs that went off patent.

So what was the problem? Why were relatively few drugs distributed in poor countries? It turns out that the problems were even more fundamental than the ability to pay. What is missing, fundamentally, are the elements of a system for delivering drugs in the poorest countries of the world. A vast shortage of doctors and medical personnel means that prescriptions for drugs go unwritten. Too few clinics are available. Pharmacies, laboratories, hospitals and medical schools are in short supply. Distribution centres don't exist. Refrigerated systems for moving vaccines and other fragile medicines are not available.

It is only by changing the underlying structure of the system that we will get people to change bad behaviour

Medical devices are difficult to repair, which means that crucial diagnostic tests may be impossible to administer. Supplies such as X-ray films, test tubes and syringes are scarce.

As a result, patent incentives don't work very well. No matter how strong the profit incentive, drug companies can't effectively get their products distributed and administered. What is needed is coordination and leadership to get the structure of the system in place so that interconnections between the elements of the system – such as incentives – can work.

So what will it take to get more doctors, nurses, physician assistants, pharmacies, laboratories, medical supplies, equipment, and everything else required to treat patients effectively in developing countries? You might think that we simply need more incentives! We need incentives for students to go to medical school and incentives for patients to go to clinics and payments to supply companies. And you would be right. All of these incentives are needed.

But this analysis would be incomplete. Incentives are only one part of a much bigger picture. A focus exclusively on incentives leads you into making a number of big assumptions about what should be done first, and about the ways people will respond to different profit opportunities. For example, imagine that you were to focus first on getting students to medical school, and imagine again that you were



successful in making progress: that students from developing countries responded to your incentive programme and enrolled on medical courses. You might feel that you had achieved a victory, and of course it would be true that the incentive had worked. And yet a problem would arise: the students that graduated from medical school under your incentive programme would go to great lengths not to practice medicine in a developing country. Why? Because there are no supplies, no medicines, no pharmacies, and so on.

Dealing exclusively with incentives is like trying to tell someone how to construct a house by discussing only the nails that will tie together the boards. A much more direct and effective way is to present the builder with a complete plan that shows what the house should look like when it is constructed, and that includes a detailed architectural diagram describing the boards, joints, and other elements of the building along with ideas about how they may be joined by nails, screws, glue and other materials.

In global health, getting incentives right is essential to effectiveness, just as nails are essential to building a house. But describing in a sequential way all the elements that should be nailed together would be enormously cumbersome for planning the construction of a house. Similarly, a plan for global health is cumbersome if all it does is deal in a piecemeal way with incentive problems one after the other. What is needed is a detailed plan.

The structure of the system

The idea that “the underlying structure of a system creates the behaviours in that system” is one of the most fundamental principles in management literature. What it implies is that individuals should not be blamed for behaving in a suboptimal way. Even if they know what is desired and even if they have incentives to do it, they still may not do it! It is only by changing the underlying structure of the system that we will get them to change their behaviours. This principle has immediate applicability in real-life situations because what it implies is that to solve any problem (business or social), we need to go beyond just manipulating the underlying incentives. Instead we must do two things: understand the underlying structure that created the problem and change this structure so as to generate behaviours that would solve the problem.

The task of changing the underlying structure of the system may sound like a tall order but as Jay Forrester commented: “All [social systems] seem to have a few sensitive influence points through which the behaviour of the system can be changed.” Find these influence

points in a system and you are halfway to achieving deep and sustainable change in that system.

What exactly is the underlying structure that needs changing? In the cases of Kitty Genovese and Angel Arce Torres, we argued that values, culture and incentives are three important elements of this structure. In the beer game, it was the time delays, the lack of information, the inability to forecast and the multiple layers in the system that formed its physical structure.

There are obviously many component parts to the structure of a system but to get to a workable definition consider the following discussion in Donella Meadows’ book *Thinking in Systems*: “A system isn’t just any old collection of things. A system is an interconnected set of elements that is coherently organised in a way that achieves something. If you look at that definition closely for a minute, you can see that a system must consist of three kinds of things: *elements*, *interconnections* and a *function* or *purpose*.”

“For example, the elements of your digestive system include teeth, enzymes, stomach and intestines. They are interrelated through the physical flow of food and through an elegant set of regulating chemical signals. The function of this system is to break down food into its basic nutrients and to transfer those nutrients into the bloodstream (another system) while discarding unusable waste.”

Thus, to understand the underlying structure of any system, we must look for three things:

1. The system’s purpose. What is it trying to achieve? This can be broken down into two finer subcategories:

- The paradigm or accepted philosophy on which the whole system is based. This is the set of unwritten and deeply-held assumptions and beliefs about how the world as a whole (and not just the system) works. Though unwritten, these beliefs are shared by all actors in the system in a subconscious way.

- The goal of the system—what are we trying to achieve? The goal is important because it determines how we organise the system and how we behave in the system (i.e. our strategy).

2. The main elements that make up the system. This can be broken down into five finer sub-categories:

- The rules in the system. These may be formal or informal, physical or social. They define what is accepted (or not) in a system and provide people with the parameters within which they have to operate.

- The physical structure of the system which includes the underlying processes as well as the structure of the hierarchy.

- Information flows in the system. What information is delivered in the system,

by whom and how.

- Time delays in the system. These may be information delays or delays between action and effect.

- Positive and negative feedback loops in the system.

3. The interconnections in the system. This can be broken down into three finer subcategories:

- The incentives in the system. These may be financial or nonfinancial, positive or negative.

- The cultural norms and values in the system. This is a set of beliefs and assumptions that acts as a moral compass for us, determining what is ‘right’ and what is ‘wrong’. It includes deeply held beliefs as to ‘the right way to do things around here,’ ‘what will get you ahead in this place,’ and ‘how this place works.’

- The way the system absorbs change (called buffers or stabilising stocks).

This might look like a frighteningly long list of things to change but we know that the task is not as onerous as it seems because of two factors. First, not all component parts of a system need to be changed for fundamental change to occur. Some component parts are more influential than others and this is where effort needs to be focused; and, second, small changes in the underlying structure of a system can produce big changes in behaviours in that system. The trick, therefore, for achieving deep and sustainable change is finding the high-leverage points in the system and focusing our efforts at achieving small changes in these parts of the system. Knowing where to focus one’s effort is half the battle won. Succeeding in introducing small changes on these points will guarantee success. ■

Costas Markides holds the Robert P Bauman Chair in Strategic Leadership and is Professor of Strategy and Entrepreneurship at London Business School. Anita McGahan is Associate Dean (Research), Director of the PhD Program and Professor of Strategic Management at the University of Toronto’s Rotman School of Management.

Resources

Ann van Ackere, Erik Reimer Larsen and John Morecroft, ‘Systems Thinking and Business Process Redesign: An Application to the Beer Game’, *European Management Journal*, Vol. 11, No. 4, 1993; Chris Bartlett and Sumantra Ghoshal, ‘Rebuilding Behavioral Context: Turn Process Reengineering into People Reengineering’, *Sloan Management Review*, 37 (1), Fall 1995; Jay W. Forrester: *Principles of Systems* (2nd edition), Productivity Press, 1968; Jay W. Forrester, *The Collected Papers of Jay W. Forrester*, Wright-Allen Press, 1975; Donella H. Meadows, *Thinking in Systems: A Primer* (edited by Diana Wright), Earthscan, 2009; Lee Ross and Richard E. Nisbett, *The Person and the Situation: Perspectives of Social Psychology*, McGraw Hill, 1991; Peter Senge, *The Fifth Discipline*, Century Business, 1990

Copyright of Business Strategy Review is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.

Copyright of London Business School Review is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.