How complex systems fail: lessons from Boeing's 737 MAX 8 crashes

A complex systems approach is a more useful way of analysing incidents such as Boeing's 737 MAX 8 failures, writes **Dr Sean Brady**

On the morning of 29 October 2018, Lion Air Flight 610 – a Boeing 737 MAX 8 aircraft – is preparing for take-off at Soekarno-Hatta International Airport, Jakarta. In command is 31-year-old Bhavye Suneja, who has more than 6000 hours of flight time, most of which were in previous versions of the 737. His co-pilot, Harvino, is ten years older, with more than 5000 flight hours.

At 6:20 am, they take off. But only minutes into the flight, Suneja's control column starts shaking. This indicates the plane is nearing a stall, a situation where the angle of the plane's wings – the plane's so-called angle of attack – is too steep, which results in a loss of aerodynamic lift. At the same time, two alerts go off in the cabin: bad altitude and air speed. Harvino asks the Captain if he wants to turn around, but Suneja says no. He asks Harvino to get clearance for a holding point to buy them some time. Harvino gets on the radio: "Flight control problem."

Then the nose of the plane suddenly dips forward. Suneja has no idea why it's happened. He presses the trim switch on his control column, which changes the angle of the small wing on the rear of the aircraft – the horizontal stabiliser. The nose of the plane comes back up. But then it suddenly dips forward again. It's like the aircraft has a mind of its own.

Beside Suneja, Harvino is working through Boeing's quick reference handbook, looking for an emergency checklist to work out what's wrong. But the handbook is no help – it says nothing about the nose repeatedly pitching downwards.

Over the next eight minutes, Suneja continues to fight with the controls. The plane repeatedly pitches forward, filling the pilots' view with the blue expanse of Jakarta Bay. And each time, Suneja flicks the trim switch, and the nose comes back up. Then it pitches downwards again. It does this 21 times, and although they are cleared for an altitude of 27,000 feet, they are still less than 6,000 feet in the air and dangerously close to the water.

Suneja asks the co-pilot to take the controls. But as Harvino takes over, the plane pitches downward again. Harvino presses the trim switch, but not as hard as Suneja had. The plane pitches further forward, then further forward again. Harvino tells the Captain they're pointing downwards. Distracted, Suneja says, "It's okay."

They plummet at 10,000 feet per minute. Harvino pulls desperately on the control column, but it has no effect. Alarms blare in the cabin: "Sink rate, sink

rate." Blue water fills their view. Harvino starts to pray. Suneja is silent. The alarms continue: "Terrain, terrain." They hit the water at an almost vertical angle, travelling at 800km/h. All 189 people on board are killed.

Over the months that followed, two narratives played out. The public narrative was driven by Boeing's CEO and top engineers, as well as by the airline regulator in the United States, the Federal Aviation Administration (FAA). As far as they were concerned, there was no fundamental issue with the 737 MAX 8. They claimed the cause of the crash was Lion Air's fault, saying that the Indonesian airline was poorly managed.

In time, the technical cause of the failure would be identified as an issue with one of the plane's angle of attack sensors. As its name suggests, this sensor measures the plane's angle of attack, which is the angle of the wing relative to the airflow. It was found to be reading an erroneously high angle, which incorrectly suggested that the aircraft was nearing a stall. But we know that aircraft are designed and built with multiple layers of redundancy, so how could an incorrect reading from a single sensor crash a plane?

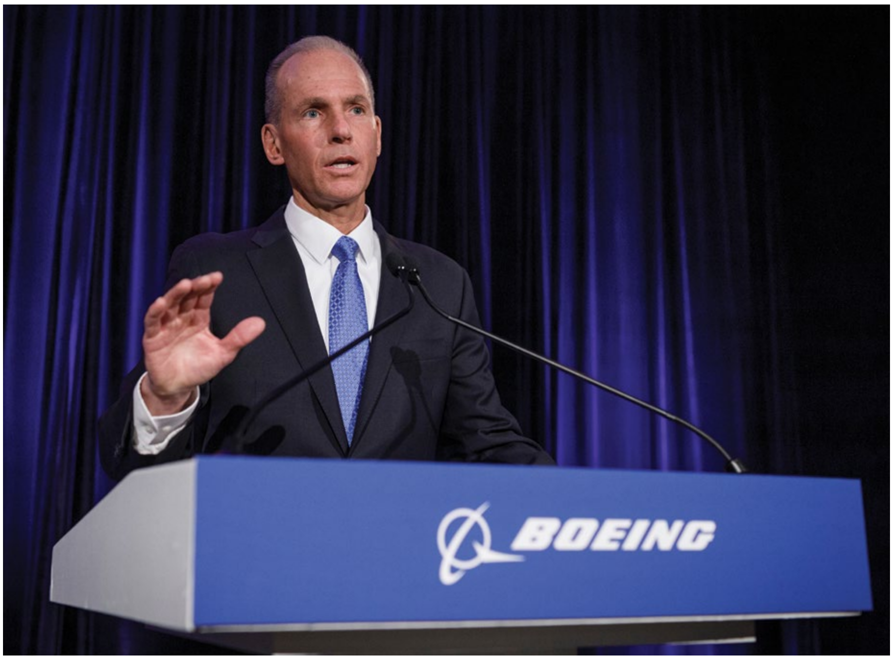

In 2019 Boeing's former CEO, Dennis Muilenburg, got the highest paycheque of his career

Complex systems

When we think about cause and effect, it would be very easy to conclude that this issue with the angle of attack sensor 'caused' the crash. In other words, the crash could have been avoided if the sensor had been working correctly. But when we examine systems as complex as the 737 MAX 8, we need to think about failure differently. We need to take a complex systems approach.

> *"In March 2019, Boeing's CEO, Dennis Muilenburg, got the highest paycheque of his career: US$31 million, including a US$13 million bonus for performance"*

Many of us think in Newtonian terms, meaning that when we examine systems, we tend to believe there is a direct link between cause and effect – everything that happens has a definite, identifiable cause and effect. Furthermore, we expect symmetry. The seriousness of the effect is related to the seriousness of the cause, significant failures happen because of significant causes, and vice versa.

However, this is different for complex systems, which are systems made up of agents or components that interact with one another and produce feedback. While we often think of systems as 'the sum of their parts', complex systems are better thought of as 'the sum of their parts and interactions'.

An analogy is a sports team. The overall performance of a sports team is so much more than the sum of the abilities of the individual players. A good team is one where the interactions between those players produce a performance that transcends the abilities of the individuals. Further, attempting to understand the overall performance of the team by studying each player in isolation will not provide much insight into the behaviour of the team as a whole.

Just as a team is the sum of its players and their interactions, in complex systems, it is these interactions that result in complex systems having a disproportionate relationship between cause and effect. This means that relatively small causes can produce very large effects. For example, the assassination of Archduke Franz Ferdinand in 1914 in Sarajevo sparked the First World War and led to millions of deaths. How could a single assassination lead to a world war?

The sand pile model

There is a different way to think about system failures, to help us on this journey: a model of a sand pile. This model, which started as a thought experiment, was developed by physicist Per Bak. As described in Mark Buchanan's *Ubiquity: Why Catastrophes Happen*, Bak and his colleagues created what physicists often describe as a 'toy model' – a model that allows you to think about complex phenomena in a simple way. Considerable research has been done on this model, and we will only examine a few key concepts here.

The model is as simple as it is profound. Bak asks us to imagine the following situation: we have a tabletop and drop grains of sand onto it at random locations, one grain at a time. As more and more grains fall, they build up into small hills. The formation of these hills is random because the grains of sand fall at random locations. As the hills grow taller, they become steeper. Eventually, one becomes so steep that an avalanche results when the next grain of sand lands on it. This avalanche could be localised, or it could trigger further avalanches as it strikes neighbouring hills.

Now, consider what is causing these avalanches. On one hand, we could say that the cause of the avalanche is the single grain of sand that fell and struck the hill (as the avalanche only occurred because this specific grain fell at this precise location). But blaming the single sand grain alone doesn't fully explain why the avalanche occurred.

Firstly, most sand grains that fall on the table do not result in avalanches. Secondly, a single grain of sand can start a small or large avalanche. The initiating event for each is the same, but the magnitude of the effect is independent of the initiating sand grain. In other words, the single grain of sand doesn't help us explain why the magnitude of some avalanches is then greater than others.

Instead, we can attribute the cause of the avalanche to the hill itself. If the hill weren't shaped as it was, the grain of sand would not have initiated the avalanche. Shifting our thinking about the cause of the avalanche, from the grain of sand to the shape of the hill, has several profound implications for understanding failure in complex systems.

If we accept that the shape of the hill – and not the sand grain – dictates the risk of an avalanche, then understanding how the hill came to be that shape is critical. We, therefore, need to know its history – how it was produced as each grain fell on it. In complex systems, we cannot take a snapshot in time, but we must consider the culmination of steps that brought us to this point.

Further, it introduces the concept of the 'critical state'. As more and more sand falls on the table and the hills get taller, the system is becoming more and more at risk of an avalanche. Bak described this process of reaching a critical state as 'self-organised criticality' because no one is organising the sand pile and increasing the risk of an avalanche. Instead, it is doing this naturally because of the interactions between the individual sand particles. As Miller & Page state in *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*: "The key driving force behind self-organised criticality is that micro-level agent behaviour tends to cause the system to self-organise and converge to critical points at which small events can have significant global impacts." The sand pile model, therefore, is a helpful way to understand why simple causes can produce significant failures in complex systems.

> *"A single grain of sand can start a small or large avalanche. The initiating event for each is the same, but the magnitude of the effect is independent of the initiating sand grain"*

For example, we can use it to re-examine the cause of the First World War. Serbian nationalist Gavrilo Princip assassinated Archduke Franz Ferdinand in Sarajevo in 1914, setting in motion a series of events that led to a world war. But when we look at this through the lens of the sand pile model, as Mark Buchanan does in his book, *Ubiquity*, he suggests that we should think of the assassination as a grain of sand and Europe as a hill in the sand pile. This hill was in a critical state due to interlocking treaties between multiple countries – a state ripe for a single grain of sand, Gavrilo Princip, to fall and start an avalanche. Once this grain landed, the interactions between the European parties cascaded and resulted in the war, just like the cascade of grains in the sand pile. If it hadn't been for Princip and the assassination, there probably would have been another initiating event. It was the critical state that mattered, not the specific grain of sand. With this in mind, let's go back and examine the story of the Boeing 737 MAX.

Boeing, Airbus and the A320neo

The story of the MAX begins not with Boeing but with its rival, Airbus - a European consortium that received its first order in the US in 1978 and who, in 1984, launched the Airbus A320 in direct competition to Boeing's existing 737. In 2010, it repeated the move with the A320neo, a plane designed to take more market share from the 737. This aircraft was larger than the previous A320 and was more fuel efficient, with the 'neo' standing for 'new engine option'. By the Paris Air Show in June 2011, Airbus had secured more than a thousand orders.

Boeing had to respond – more than a third of their profits came from the 737 – but they didn't have an aircraft that could compete with the fuel efficiency of these new planes. Then in July that year, they got word of a potential deal between American Airlines and Airbus. It looked like the airline was about to order the neo. Boeing stepped in to try and secure its own deal, convincing American Airlines to split the order: the airline would buy 260 A320neos, with the remaining 200 planes being a more fuel-efficient aircraft from Boeing. This plane, which was entirely hypothetical at this point, would come to be named the 737 MAX 8.

Developing the 737 MAX 8

The original Boeing 737 was launched in January 1967. By 1988, it was flown by over 137 operators worldwide and described as the 'unsung prodigy' of the Boeing family. It would also form the basis of the 737 MAX 8. Boeing planned to take the existing plane, replace the engines with more fuel-efficient ones, and bring them forward on the wings.

A critical decision made early in the MAX's development was that Boeing wanted pilots already trained on the existing 737 to be able to fly the MAX with no additional simulator training. Training is a significant cost for airlines: simulators cost around US$15 million each, pilots have to be taken out of service, and the training itself costs hundreds of dollars per hour. In fact, training, wages, and maintenance costs amount to 20 per cent of the overall costs of running an airline – more than they spend on fuel. If Boeing could put a new plane on the market without requiring pilots to undertake additional simulator training, it would give them a massive advantage. But to achieve this, they had to ensure they could modify the existing 737 and not add any new functionality that would change the handling or operation of the aircraft.

To understand the environment in which the development of this aircraft occurred, we need to look at Boeing's history and the dramatic changes it went through from 1997 onwards. The company was founded in 1916. By 1944, it had a workforce of 50,000 people, and by the 1960s, this had jumped to 142,400. They were all about producing high-quality, safe planes and had a saying: "We hire engineers and other people". At meetings, designers were encouraged to fight loudly for what they wanted on the planes to make them safer.

Dr Sean Brady says the sand pile model is a helpful way to understand why simple causes can produce significant failures in complex systems

But in 1997, they merged with McDonnell Douglas – a company much more cut-throat when it came to cost-cutting. As the McDonnell Douglas executives spread throughout the organisation, their approach to building planes began to dominate how Boeing operated. The infiltration was described as 'hunter killer assassins' being let loose on a room full of engineers. The organisation started to change – in many ways, a microcosm of the Jack Welch-inspired culture of the times. Cost-cutting and a return on shareholder investment seemed more important than producing quality aircraft. Engineering views took a back seat. The 'other people' were now firmly in charge.

Against this backdrop, the design of the 737 MAX 8 got underway, with a focus on 'more for less'. A countdown clock was set up in the conference room where program meetings took place to remind people there was no time to waste. Overshadowing every decision was the drive to ensure no new functionality was added to the aircraft, which would have required additional pilot training.

The Manoeuvring Characteristics Augmentation System

A big problem with Boeing's design approach for the MAX emerged during wind tunnel tests on a scale model of the aircraft. The model pitched up during tight, high-speed turns due to the new engines being placed further forward on the wings. This behaviour was a genuine concern: if the plane pitched up too far, its angle of attack would become too steep and it would stall, which could lead to a crash.

The 737 chief pilot, Ray Craig, examined the problem and discovered it only happened in the part of the flight envelope that commercial pilots rarely go. But pilots could enter this zone if they were dealing with high turbulence or responding to some upset. And if they did, the nose of the plane could pitch up and they could stall.

This issue had to be addressed and several mechanical solutions were proposed. They explored putting tiny vanes on the wing, but they didn't think that would work. The only real mechanical solution was redesigning the tail and removing the pitch-up risk. But this was a costly solution that could delay the plane's release.

So they agreed upon a software, not mechanical, solution. This software system went by the cumbersome name 'Manoeuvring Characteristics Augmentation System', or MCAS for short. The software would detect when the plane was pitching up too far while in this edge-of-the-envelope zone, and it would rotate the horizontal stabiliser at the back of the aircraft and push the plane's nose back down. This would manage the stall risk. Not only was this solution cheap, but it would ensure the plane handled like the previous 737s – simulator training would not be required.

To detect when the plane was pitching upwards, the software would rely on measurements from two sensors: an accelerometer measuring the plane's acceleration, and one angle of attack vane mounted on the front of the aircraft, measuring the plane's angle of attack. While there are two angle of attack vanes on the 737 MAX 8, the software would use only one. Critically, the software would rely on input from two sensors, not one, to control the plane in this edge-of-the-envelope zone. The chief pilot of the project, however, wanted a hardware solution, but was overruled because the software was cheaper.

But adopting this solution did raise a concerning problem. Boeing engineers were worried about what they should call this software and who they should tell about it. There was a real danger that the regulator, the FAA, might view this software as 'new functionality'. And if they did, it was something that pilots would need to be trained on in the simulator. This was the last thing Boeing wanted: they had publicly announced that existing 737 pilots could migrate to the MAX by undertaking a short training session on an iPad.

The extension of MCAS

The design and development of the aircraft continued, with the first test flight of the 737 MAX taking place in January 2016. But there was more bad news only a few months into these tests. The pitch-up problem, which the scale model showed only happened near the edge of the envelope, was now also evident at slower speeds. This meant that an edge-of-the-envelope concern could now occur during routine operations. To make matters worse, this stall risk at slower speeds could occur during take-off and landing - the most vulnerable part of the flight, and when the pilots are at their busiest. There was now a genuine concern within Boeing that the FAA would not certify the plane.

> ***"Pilots could now take the plane up without additional simulator training, with a system on board that was not discussed in the manual, which could override their control"***

To solve this problem, Boeing extended the software solution to cover these low-speed stall risks. If the aircraft was at risk of pitching up at these slow speeds, MCAS would detect it and activate the horizontal stabiliser on the tail so the plane would pitch down again. This stabiliser, which MCAS could move by 0.6 of a degree previously, could now move 2.5 degrees at slower speeds. But at these slower speeds, the software could no longer use the accelerometer as an input. MCAS now relied on only one sensor: an angle of attack vane. Boeing, who had previously put so much focus on engineering and safety, were now relying on a system with no redundancy if anything was to happen to this sensor.

And Boeing had another problem. They'd picked an unproven supplier to deliver the simulator. While Boeing wanted no simulator training for existing pilots, they still needed a simulator for new pilots who had never flown a 737. But simulator development was falling behind. This not only proved a worry for training new pilots, but it also meant that if the FAA declared that training was required for the MAX, even for pilots who had flown the 737 before, there was nowhere for this training to take place. That would prohibit planes from flying. Now the necessity of convincing the FAA there was no new functionality, specifically around the role that MCAS would play, was crucial.

MCAS was first loaded onto the 737 MAX 8 computer on 15 August 2016 – it was now ready for production. In November of that year, Boeing engineers sent their system safety assessment of MCAS to the FAA. The latest version, revision E, had all the details of how the system would operate at lower speeds and move 2.5 degrees. But revision E was not the version submitted to the FAA. Instead, revision C was submitted, covering only MCAS's more limited role. Not only did the FAA approve MCAS, but they also approved no reference being made to it in the manual.

Pilots could now take the plane up without additional simulator training, with a system on board that was not discussed in the manual, which could override their control. Most MAX pilots didn't even know MCAS existed.

The first crash

The grain of sand that would initiate the Lion Air 610 crash was a misaligned angle of attack sensor that erroneously told MCAS that the plane's angle of attack was too steep. MCAS engaged, activated the horizontal stabiliser, and pushed the nose of the aircraft down. While Captain Suneja could activate the trim switch and pull the nose back up, each time he did, the software would continue receiving data from the misaligned sensor, reactivate, and push the nose back down again. Even when Harvino pulled back on the control column to pull the nose up, this had no effect – MCAS was designed to override it. As Suneja and Harvino continued to battle the aircraft's behaviour, neither was aware of the software, what it was doing, or what was required to deactivate it. They were entirely at its mercy.

In November 2018, one month after the Lion Air crash – with planes still flying – Boeing met with pilots and trainers. The pilots and trainers were shocked when they heard about MCAS and the fact that they had not been told it was on the plane. Boeing also explained what was required to disable MCAS in the event of a malfunctioning angle of attack sensor. This sequence would turn out to be very difficult to execute in the real world.

Also that month, the Indonesian investigators released their report on the Lion

Air crash, primarily blaming the pilots and maintenance staff. The MAX continued to fly, Boeing's stock price rose over the following months, and the FAA gave them ten months to fix the software, even when the FAA's own analysis concluded that the MAX posed a serious risk. In March 2019, Boeing's CEO, Dennis Muilenburg, got the highest paycheque of his career: US$31 million, including a US$13 million bonus for performance.

The second crash

And then, on 10 March 2019, only five months after the Lion Air crash, Boeing received news of a second incident. Ethiopian Airlines flight 302, a 737 MAX 8, had taken off from Addis Ababa Bole Airport. At the time, Ethiopia Airways were considered one of the best-run airlines in Africa. But shortly after take-off, MCAS activated because the angle of attack sensor developed an electrical issue. It was sending incorrect data to the software. The pilots fought against MCAS, trying to execute the sequence Boeing had prescribed for disabling the software. But this was a complex and ill-explained sequence that proved very difficult to execute in flight. The aircraft crashed, tearing itself apart and killing all 157 onboard.

By now, there had been two crashes and 346 people had been killed, yet Boeing still publicly argued the MAX 8 was okay – and the FAA agreed. China moved first and grounded the plane. They were followed by the European Union, India, Australia, Singapore, and Canada. The US grounded it on 13 March 2019.

Seeking closure

A traditional approach to understanding Boeing's 737 MAX 8 failures would result in us attempting to draw a line between cause and effect, beginning with the issues with the angle of attack sensors. After all, if it weren't for these faulty sensors, the crashes would have been avoided. But taking a complex systems approach, especially through the lens of the sand pile model, provides a more useful way of viewing these types of incidents.

Rather than trying to string all the contributing factors together in a line, the sand pile model asks us to consider how each of them interacted with one another and layered upon one another to build a hill. It asks us to examine the change in culture at Boeing from engineering excellence to cost cutting; the need to get a new aircraft out quickly and cheaply in order to compete with Airbus; the decision on no simulator training for existing 737 pilots; the use of MCAS, and then the extension of that use; the software's reliance on a single sensor; the fear the FAA wouldn't certify the aircraft if MCAS was deemed new functionality; and the decision not to tell the trainers and pilots about the software, nor provide details of it in the manual. It asks us to treat the issues with the angle of attack sensors as the initiating event, with the failure being the result of the shape of the hill we've built, not the grain of sand we've dropped.

It asks us to re-examine our more traditional views on cause and effect, and instead look more closely at the sand piles we build in our own projects and organisations. It requires us to ask ourselves if the systems we've built are tending towards a critical state, just waiting for that single, innocuous grain of sand to bring them tumbling down.

Dr Sean Brady is a forensic engineer and the managing director of Brady Heywood, where he works with businesses, governments, and the legal sector to investigate and resolve complex issues that typically require a systems approach. He has acted as an expert witness in numerous proceedings involving a wide range of constructed facilities, and authored the Brady Review into mining fatalities that was tabled in the Queensland Parliament in 2020.