

1 “BRACE FOR IMPACT”

One essential characteristic of modern life is that we all depend on systems – on assemblages of people or technologies or both – and among our most profound difficulties is making them work.

– [Atul Gawande](#), 2010, *The Checklist Manifesto: How to Get Things Right*

On a cold January afternoon in 2009, US Airways Flight 1549 departed LaGuardia airport in New York, headed for Charlotte, North Carolina. There had been a light dusting of snow that morning but even with cloudy skies the scenery was beautiful enough along the Hudson River for Captain Chelsey “Sully” Sullenberger to remark, “What a view of the Hudson today!”¹

Less than two minutes after takeoff, still rising and heavy with fuel, the plane hit a flock of Canadian geese. Sullenberger and his co-pilot, Jeffrey Skiles, saw them coming, but there was no time to avoid. “Birds,” said Sullenberger. “Whoa,” said Skiles, “oh, [explicative].” Even a single chicken will destroy an engine. In a moment, multiple large geese were sucked in, crippling the plane. Communication and other systems remained, but the thrust pushing the plane into the air was gone, leaving the crew piloting a fast-descending 150,000 pound glider. “Uh oh,” said Skiles.

Sullenberger and Skiles had to take action quickly. Sullenberger contacted TRACON, the system of air traffic controllers that provides assistance in areas with multiple airports. “Hit birds, we lost thrust in both engines. We’re turning back toward LaGuardia.”

Air traffic controller Patrick Harten confirmed this decision and started preparing an emergency flight path and clearing runways. It's clear from the transcripts that loss of all engines was unprecedented. "He lost all engines," Harten said. "He lost the thrust in the engines. He is returning immediately." The reply from the airport tower was, with a tone of disbelief, "Which engines?!" "He lost thrust in *both* engines, he said," Harten responded. "Got it," the tower controller replied, and initiated emergency procedures.

It was only thirty-six seconds after announcing the bird strike that Captain Sullenberger gave up returning to LaGuardia. "We're unable. We may end up in the Hudson." The air traffic controller didn't immediately change the plan, and offered an open runway at LaGuardia. Sullenberger replied with one word: "Unable."

The air traffic controller acknowledged but offered yet another runway at LaGuardia. Sullenberger responded, "I'm not sure if we can make any runway. Oh, what's over to our right? Anything in New Jersey? Maybe Teterboro?" The air traffic controller contacted the Teterboro airport and confirmed an open runway to Sullenberger, all in less than sixteen seconds. During that time, though, Captain Sullenberger had assessed the situation and made his decision. "We can't do it," he said, exactly two minutes since the first emergency communication. "We're gonna be in the Hudson." He turned on the intercom to broadcast to the cabin.

"Brace for impact," he said.

Miracle or ...

Sullenberger's water landing and rescue of all 155 people on board has been called the "Miracle on the Hudson." But was it a miracle, or was it the product of decades of engineering and design choices, training regulations, and semi-autonomous systems incorporated into the brain of the plane itself? Certainly, luck played a role: The weather was clear. The accident occurred in the daytime. The Hudson River was nearby, fairly clear for a landing, and minutes from rescue boats. A cold front brought an unusual placidity to the Hudson, but without large chunks of ice in the water.

But it wasn't all luck. Many contributors to the Miracle were by design and under human control: Captain Sullenberger was highly experienced, with more than 20,000 hours of flight time since starting

his commercial career in 1980. Before that, he was a fighter pilot in the US Air Force. He was also an expert in aviation safety, with his own consulting business and advanced degrees in Industrial Psychology from Purdue. First Officer Skiles also had over 20,000 hours of logged flight time and had flown at the rank of captain himself.² Although no runways ended up being used, Harden and the airport controllers communicated quickly and effectively. Everyone in the situation was consulting instrument panels, predictive systems, checklists, and displays throughout, each carefully designed to provide information and support fast decisions.

Thus, Sullenberger and Skiles landed safely because of their own judgments combined with the support systems engineered into the airplane and into the procedures they followed. Steven Johnson summarized the contributions of automation and decision aids poetically in his book *Future Perfect*:

Most non-pilots think of modern planes as possessing two primary modes: “autopilot,” during which the computers are effectively flying the plane, and “manual,” during which humans are in charge. But fly-by-wire is a more subtle innovation. Sullenberger was in command of the aircraft as he steered it toward the Hudson, but the fly-by-wire system was silently working alongside him throughout, setting the boundaries or optimal targets for his actions. That extraordinary landing was a kind of duet between a single human being at the helm of the aircraft and the embedded knowledge of the thousands of human beings that had collaborated over the years to build the Airbus A320’s fly-by-wire technology.³

However, an in-depth look into the Miracle also revealed problems with the flight systems. As we identify these problems, the landing on the Hudson becomes a perfect microcosm for how to make improvements in the human factors of future flights. The first order of business is to find the gaps between what pilots are capable of and what is demanded of them, and then what systems are or could be in place to close those gaps.

In the aftermath of the Hudson landing, Sullenberger and Skiles were scrutinized for their choices. Flight simulators were set up by the National Transportation and Safety Board (NTSB) and pilots tried to land on the Hudson or were asked to return to various airports. Some of

these airport landings were successful, prompting gleeful headlines such as "Sully Could Have Made It Back to LaGuardia."⁴ However, the simulator pilots knew what they would face ahead of time and reacted instantly, even then not always succeeding in landing at an airport. Only one attempt was made with a delay added after the birds were hit to simulate human decision time. When a delay was added, the pilot did not make it to an airport in the simulation – all onboard would have died.⁵

We tend to respond to critiques of our heroes by pushing back, or considering any questioning of them to be an insult. If the questions are meant to identify scapegoats and assign blame, then we are right to be upset when our heroes are questioned. Blaming "pilot error" on a scapegoat does little to prevent future incidents. But if we search for continuous improvement, then it is always correct to question those heroes to analyze what went wrong, what went right, and how we can make our systems even better. For example, the NTSB report noted that few pilots were able to hit the water at a good angle in the simulator, but that the one who did used a specific technique: "approaching the water at a high speed, leveling the airplane a few feet above the water with the help of the radar altimeter, and then bleeding off airspeed in ground effect until the airplane settled into the water." Thus, this new technique was learned from investigating the Hudson crash and is now taught to other pilots. Other findings by the NTSB were that there was no checklist for ditching a plane at low altitude – an issue the FAA then addressed. Trying to learn and improve is a core principle of good design, but one that is not as natural as the urge to blame. As an educated public, we must also insist that industries and governments adhere to the principle of learning and improving when creating products, regulations, or meting out punishments to the "bad apples."

"Caution! Terrain! Pull Up Pull Up!"

One of the devices designed to intervene in or before an air emergency is the Enhanced Ground Proximity Warning System (EGPWS). When crewmembers are distracted or have to attend to other issues, the EGPWS checks the aircraft position relative to the ground and obstacles. If the aircraft changes altitude too quickly, the system will verbally instruct "Don't sink!" If the aircraft is too close to the ground, particularly mountains, the system will start with information

“Too low, terrain,” then instruct “Pull up, pull up” along with the reason why, “Terrain,” and an alert word (“Caution!”) to help the pilot know action must be taken. Terrain warning systems may be the biggest advancement in aviation safety since the 1970s – after their introduction, accidents involving “flying a perfectly good plane into the ground” dropped dramatically, almost to nonexistence.⁶

However, as with all human-engineered systems, the EGPWS is not always reliable. Most frequently it fails by offering false alarms in a loud distracting voice. If an airport is not programmed into the EGPWS (many are not) or the plane is making an emergency landing at a non-airport, then it will alarm constantly during the descent. Pilots are annoyed at the non-stop sound when landing at rural airports – and it is very very loud. The EGPWS can be heard in the cockpit audio from the Hudson landing because the system did not know the Hudson landing was intentional, adding more stress to an already stressful situation.

- EGPWS: Too low, terrain. Too low, terrain. Too low, terrain. Caution, terrain. Caution, terrain. Too low, terrain. Too low, gear.
- SKILES: Hundred and fifty knots. Got flaps two, you want more?
- SULLENBERGER: No let's stay at two. Got any ideas?
- HARTEN: Cactus fifteen twenty nine if you can uh . . . you got uh runway uh two nine available at Newark it'll be two o'clock and seven miles.
- EGPWS: Caution, terrain. Caution, terrain.
- SKILES: Actually, not.
- EGPWS: Terrain terrain. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. Pull up. (Repeats indefinitely in background)
- SULLENBERGER: We're gonna brace.

As someone who has to turn down the radio to be able to concentrate when merging onto the highway, I felt for these pilots.

In a 2012 crash into Mount Salak, Indonesia, the crew had disabled the EGWPS system when they believed it was malfunctioning. All forty-five passengers died.⁷ In a 2010 flight from Poland to Russia, with the president of Poland onboard, the Russian airport was not

programmed into the EGWPS. The system sounded an alert as the plane approached, but it was ignored because the crew knew the airport was not programmed and they expected the alarms. Unfortunately, the alert was really about the trees and terrain they were going to hit before reaching the runway, rather than the "normal" alarms that they were landing at the airport. No one survived.⁸ These accidents tragically illustrated the balance between human trust in the EGWPS automation and its reliability. When the systems are not trusted, they are turned off or ignored. At other times, the system may behave as designed, going off as altitude declines, but is a distraction from the emergency at hand. Finding the right balance depends on advance testing of these systems, because even the best automated system *will* fail. We can learn from these events for future designs, such as autonomous cars or drones. How failures occur, and how they affect the person in the vehicle or around it, is within our control. Will it fail gracefully, with back-up systems ameliorating the danger of the failure? Will it be transparent in its failure, so that the operator, pilot, or driver understands what is failing and when? Will it be obvious how to react to the failure, quickly and accurately? Answering these questions during design will give us the best chance at avoiding future tragedies.

Other decision aids in the cockpit included the electronics that partially automated flying the plane. The plane calculated and displayed the best gliding speed for Sullenberger and held itself to that speed, also displaying how it anticipated changing speed ten seconds in the future. This freed Sullenberger to focus on other decisions, rather than having to hold the plane to the right speed. Human factors psychologists and engineers call picking a person or machine for a job *function allocation*: gliding speed was allocated to the machine.

Function allocation means to consider what jobs best suit machines and what jobs best suit humans. For example, human reaction time is slow compared to computers, so when a fast action needs to be taken, especially if there is a reliable cue that prompts that action, it should be allocated to a machine. Harkening back to the EGWPS, checking an altitude boundary and issuing the command "Pull up" below that boundary is a simple job for a machine. Forcing a human to remember to check altitude while multitasking or in an emergency demands a great deal of time and attention. However, when it comes to visual pattern matching (e.g., picking out a target) or making decisions

based on multiple ambiguous cues (e.g., diplomatic negotiation), humans have the edge. But making those kinds of decisions can be effortful, so it's often best to allocate as much of the lower-level decisions to a machine to free up the resources of the human for those tough questions.

Checklists Help with Decisions

Sullenberger and Skiles immediately went to the reignition checklist when their engines were destroyed and continued following it in tandem with water landing procedures, only stopping once the plane was in the Hudson. Following checklists is ingrained in aviation and is slowly becoming standard in other domains, such as health and medicine. However, even with the best checklist, humans may need to decide what and when to follow. When testifying in front of the National Transportation Safety Board, Sullenberger said,

We didn't have time to consult all the written guidance, we didn't have time to complete the appropriate checklists. So Jeff Skiles and I had to work almost intuitively in a very close-knit fashion without having a chance to verbalize every decision, every part of the situation. By observing each others' actions and hearing our transmissions and our words to others, we were able to quickly be on the same page, know what needed to be done, and begin to do it.⁹

They also needed to prioritize their actions. In an interview with *Air and Space Magazine*, Sullenberger said,

The higher priority procedure to follow was for the loss of both engines. The ditching [landing outside an airport] would have been far secondary to that. Not only did we not have time to go through a ditching checklist, we didn't have time to even finish the checklist for loss of thrust in both engines. That was a three-page checklist, and we didn't even have time to finish the first page. That's how time-compressed this was.¹⁰

Sullenberger's decision to ignore the three-page checklist, meant to be used at 30,000 feet instead of 3,000 feet, exemplifies the importance of usable design. In an interview with Jon Stewart on *The Daily Show*, the taciturn Sullenberger acknowledged how poor design

reduced the information he could gather as the plane went down. Stewart commented with disbelief,

You said your partner reaches over for the manual. I guess there's a manual they put in the cockpit for things that go wrong, and it used to have tabs on it for easy [access] – like, "blue tab is plane going down." And as a cost-saving measure they had removed the tabs. So he was literally going like, "I better check the index!" You know? Like, how crazy is that?

Sullenberger responded: "It's one of those minor things that by itself might not make a big difference. But, you know part of what we do is manage risk. We look for ways to make the system better. And I think that would make the system better, if we put the tabs back on."¹¹ Finding instruction quickly (in Sullenberger's case more quickly than anyone had imagined) could have been easier with search tabs, or a fast electronic search, or a just-in-time display fed by artificial intelligence. The ways to support are as unlimited as human imagination paired with engineering – but all of them should be tested for ease of use, especially in time-critical emergencies with other alarms and systems going off.

Checklists Put Everyone on the Same Page

Checklists ensure an entire team understands the past, present, and anticipated future of their job. This shared mental state requires *theory of mind*, a term first coined by developmental psychologists to describe how small children move from having their own thoughts to understanding that other people *also* have thoughts, and those thoughts can be different from theirs. It is hard to imagine now, but when we were small children we did not know that the people around us could be thinking, knowing, or seeing things differently than we do. Incidentally, this also means small children are terrible (even incapable) liars, since they believe you already know everything they are thinking. It also explains their frustration when you don't seem to be able to read their minds to know what it is that they want.¹²

Theory of mind is not just for kids. It persists in adulthood in small ways. Adults understand other people have their own thoughts and experiences, but we still tend to believe others think more like we do than is true and are often shocked when confronted with just how differently another person thinks. Just read the comments section on

any news site. Research studies on adult theory of mind often have a person try to communicate an idea to another (one did this through drumming out a song on a table) and have them judge how well they believed the other person understood their intent. Then, the researchers compare that judgement to the other person's actual understanding. In the drumming study, people overwhelmingly thought the other person would "get" the song – it seemed so obvious to the drummer, how could it not be obvious to the receiver? But it was not obvious. Hardly any receiver guessed the correct song from hearing the drumming (the exception being "Jingle Bells").¹³ This one you can try at home.

Checklists in medicine can help prevent theory of mind mistakes from impacting surgical outcomes. I worked with a team of veterinary cardiologists in 2015 to develop a surgical checklist for their procedures and one important step in the checklist was for the surgeon to review "anticipated critical events and unexpected steps" and "expected operative duration" with the rest of the team. This was because the surgeon fully understood the medical history of the animal, including age and other potentially complicating variables like body size or severity of the heart problem. It would be easy for the surgeon to think the whole team shared in all of this knowledge and expectation, but there was a wide variety of knowledge, experience, and roles distributed among the surgical technicians, anesthesiologist, and anesthesia technicians. An explicit call for communication in the checklist made sure everyone was working with the same information, had similar expectations, and could make better decisions based on that information.¹⁴

Thus, human decisions are aided in two ways by a checklist. Checklists support fallible human memory and encourage explicit communication within a team. The quote from Dr. Atul Gawande that began this chapter captured the messages of our failures, as one step or job may be simple enough, but complex systems such as aviation and healthcare overwhelm the human brain.¹⁵ It is the formal systems we create, such as checklists, that address the "profound difficulties" of making them work.

Conclusion

The secret to the survival of our species is our adaptability. We have created an impossible technological world, one where we shouldn't be able to function, where machines carry us too high for

oxygen in the atmosphere and too fast for our reaction times. Yet, we have also created systems to support our capabilities and overcome our limitations. This means that our heroes are not superheroes, nor do they need to be. They are human. The intense training for pilots, surgeons, and police is important, but it is augmented by the technology that supports their decisions and actions. Acknowledging that we're not perfect, that we need support from checklists or automated systems, is the first step to being able to accomplish the extraordinary, such as keeping millions of flights safe in the air each day.

The next step is to make sure those support systems and automations are designed to fit with our all-too-human limits: checklists can't be too long, automation needs to account for failures and false alarms, and we need repeated reminders that other people think differently than we do. We need to keep an eye on the technology – if it's poorly designed, it can do more harm than good. We must insist on, and dedicate resources to, a culture of improvement, where we analyze what goes wrong but also what goes right. The Miracle on the Hudson illustrated every aspect of such a culture, from the well-tested and helpful fly-by-wire systems on the plane, to the after-incident investigation and tests, which improved training and interfaces for future flights.