Failure

Why Science Is So Successful

Stuart Firestein

OXFORD
UNIVERSITY PRESS

Contents

Introduction

Perhaps the history of errors of mankind, all things considered, is more valuable and interesting than that of their discoveries. Truth is uniform and narrow . . . but error is endlessly diversified.
—Benjamin Franklin

This book has failure written all over it.

Literally of course, but metaphorically as well. So failure will stalk this book, and it may occasionally win a round. But if I get it right, you'll understand that those failures are an important part of the book, an absolutely necessary ingredient. A book on failure cannot just be a lecture; it must be a sort of demonstration as well. And so now, by some sleight of hand, I have at least partly inoculated myself against failing by telling you that the theme of the book is how important failures are. Come to think of it, that's also a theme—that we must make and defend a space for noncatastrophic failure, a place where failure can happen regularly.

This book follows, and is kind of an extension, of another book I wrote recently called *Ignorance: How It Drives Science*. As you can see I'm carving out a nice, neat, little niche for myself. It might seem like I am becoming a merchant of despair. In fact I find both of these subjects to be uplifting. Although ignorance and failure are commonly thought of in a negative light, in science they are just the opposite: they are where all the interesting action is. This will be a key point in this book—that failure in science is fundamentally different from all the other failures you've read about in self-help and business books and articles in *Wired* and *Slate*. It is a kind of failure we don't appreciate enough. Not understanding this, not appreciating failure sufficiently, leads to distorted views of science and denies one a surprisingly useful but rarely considered version of failure. I hope that on this one point I don't fail you.

Science, the great intellectual achievement of modern Western culture, is often depicted as resting on pillars of great foundational strength and intellectual might.

These pillars are variously identified as KNOWLEDGE and REASON, or FACT and TRUTH, or EXPERIMENT and OBJECTIVITY. Quite impressive. Students are regularly asked to approach science in the reverential way that these ponderous pillars demand. Perhaps such pillars are the correct depiction for textbook science—the stuff that is frozen in time and that generations of these same poor students have been required to master, by which we usually mean *temporarily memorize*. But then there is current science, the real stuff that goes on every day in laboratories and in minds across the world. Science rests, I am afraid to say, on two somewhat less imposing sounding pillars—IGNORANCE and FAILURE.

Yes, that's it, the whole tremendous edifice. The costly research programs, the years of education, the dedication of cadres of PhDs, teetering on top of Ignorance and Failure. But without these two, the entire business would come to a standstill. In fact, Ignorance and Failure are not so much pillars as engines that propel science forward. They make it at once a reckless endeavor and a conservative process, a creative enterprise composed of mind numbing reams of data. I understand that this view of science, beholden to Ignorance and Failure, is probably not the common perception, and that few outside of practicing scientists will immediately recognize the truth of this proposition. But I bet that anyone who has made a career of science, reading this now, is nodding in agreement. Indeed, every scientist to whom I have mentioned that I am writing a book on failure has immediately offered to contribute a chapter! Remarkably, most of us make a pretty good living doing this kind of work and virtually every scientist I know loves their work. So how could that be, composed as it is mostly of ignorance and failure—with perhaps a dash of accident or serendipity thrown in?

It may seem that I'm putting you on here, pretending to reveal some dirty little secret just to grab your attention. But the thing is, it's not a secret at all: it's general knowledge, inside of science. Somehow outside of the scientific establishment it seems we do a very bad job of letting everyone else in on what we do. So many things just get taken for granted that it never occurs to us to make it explicit. You know more or less what lawyers do, what accountants do, what journalists do, what car mechanics do—even if you couldn't do any of those things yourself. But when I tell a crowd of my daughter's friend's parents that I'm a scientist, all they want to know is what I do. Actually *do*, during the day, every day.

One curious thing about this book was that it never coalesced into a linear argument with some internal logic driving it forward. I didn't start the chapters in any particular order, and I kept working on them in no particular order. They are more like essays than chapters, each a reflection on some aspect of failure and science. The famed immunologist and science writer Sir Peter Medawar wrote a piece for the *Saturday Review* titled "Is the Scientific Paper Fraudulent?" His claim was not that scientific papers were untrue, but that they were constructed in a way that did not reflect the actual experimental or intellectual processes at work. They were reconstructed in some narrative order intended to drive home the point but were not an accurate record of the way it actually happened. This book is sort of just the opposite. It has not been put together in some carefully logical order that builds to a convincing and unassailable argument. It's more a collection of ideas, some of which I hope are new to you. They were to me.

One of the things I do hope this book will accomplish is to show science as less of an edifice built on great and imponderable pillars and more as a quite normal human activity. I don't mean by knocking it down a peg or two, but rather by building it up as a remarkable and surprisingly accessible way of seeing the world. Science is accessible to everyone because really, at its core, it is all about ignorance and failure, and perhaps the occasional lucky accident. We can all appreciate that.

ONE

Failing to Define Failure

A real failure does not need an excuse. It is an end in itself.
—Gertrude Stein

I have chosen this deceptively simple sounding statement, so typical of Gertrude Stein, to open this book because it gets so quickly to the heart of the matter. It challenges, right from the beginning, our idea of what a failure may be. What kind of a failure is Stein talking about here? What makes a "real" failure? Are there "unreal" failures, or lesser failures?

Like so many important words, *failure* is much too simple for the class of things it represents. Failure comes in many flavors, and strengths, and contexts, and values, and innumerable other variables. Nothing just stands alone as a failure without knowing something more about it. In the famous *Encyclopédie* of the French enlightenment, Diderot and d'Alembert (1751–1772) under the entry for *erreur*, which seems intended to cover failure as well, caution that there is no way to develop a general description or classification because *erreur* comes in so many forms. I started this project with what I thought were a few clear ideas about failure and its value in the pursuit of scientific explanations. What surprised me was how quickly those few ideas generated dozens of questions.

There is a continuum of failure, not just one narrow kind. Yes, there are failures that are just mistakes or errors, and they may often be no more than an unfortunate waste of time. There are failures from which you learn simple lessons: be more careful, take more time, check your answers. There are failures that can be taken as much larger life lessons: a failed marriage, a failed business venture; painful but perhaps character building. There are failures that lead to unexpected and otherwise unavailable discoveries: they often seem like serendipity, an accidental failure that opened a door you didn't even know was there. There are failures that are informative: it doesn't work this way; there must be some other way. There are failures that lead to other failures that eventually lead to some kind of success about learning why the other paths were failures. There are failures that are good for a while and then not—in science you might think of alchemy, a failure that nonetheless provided the foundations of modern chemistry.

Failures can be minimal and easily dismissed; they can be catastrophic and harmful. There are failures that should be encouraged and others that should be discouraged.

The list could go on. But I don't want to get sidetracked into a lengthy polemic trying to define failure, which would surely fail. We'll come upon all sorts of failures as we proceed, and we would do best to think of them as discoveries, not contradictions. Rather, I want to focus on the *role* that failure, in all its many identities, plays in science and how it contributes to making it such a successful enterprise.

Stein seems to be complaining about the common response to a failure—which is apology. Failure as mistake, unintended or unavoidable or because of some shortcoming that you are responsible for. Failure as the result of stupidity and naiveté that requires excuses and apologies. Why did you let that fail? Can't you do any better than that? Or, perhaps less antagonistic but no less disappointing, failure as inevitable. Well, that wasn't likely going to work. What did you expect? What a stupid thing to have even tried. And so forth. Stein, in that first simple sentence, identifies all these bad failures, useless failure, failures that demean failure.

Instead, how about failure that stems not from ineptitude, inattention, or incapacity. (True, even those occasionally turn out to reveal something unexpected and sometimes wonderful. But I wouldn't depend on them. Sloppy indifference can get you only so far.) A real failure is different from all those that need or are accompanied by an excuse—because it needs no excuse.

So what are good failures? Ones that need no excuse and are an end in themselves? Not really an end in the typical sense—that is, not an end where you give up trying anything else. Rather an *end* in the sense of something new and valuable. Something to be proud of and therefore requiring no excuse, even if it was "wrong."

Are there really such failures? Of course there are the mistakes we learn from, the errors that can be corrected, the failures that can be turned to success. But I'd like to take a chance here and venture that Stein meant something deeper than that. That she really meant meaningful failure. In the limit, this could mean that you might produce nothing but meaningful failures for your entire life and still be counted a success. Or at least never need to apologize. Is that really possible? What are these magical failures?

I have two possible answers. The first is that failures that are ends in themselves are interesting. Interesting is another word that one has to be careful about. It's easy to use, but then it's kind of vague and subjective. Is there anything that's interesting to everybody? I doubt that. But if we take interesting as a descriptor rather than an identifier—that is, a quality of something and not necessarily a particular thing itself—then we can perhaps come to an understanding. When the same Gertrude Stein was asked to write a piece about the atom bomb (shortly after its use in WW2 and, as it turns out, shortly before her death in 1946), she responded that it held no interest for her. She liked detective stories and that sort of literature, but death rays and super weapons were not that interesting because they left nothing behind. Someone sets off a bomb or some weapon of mass destruction that kills everybody and ends everything. So what's to be interested in? Certainly better if it didn't happen, but if nothing is all you're left with, then who cares? So maybe it's what's left that could make something an interesting failure. Good failures, we could call them Stein Failures, are those that leave a wake of interesting stuff behind: ideas, questions, paradoxes, enigmas, contradictions—you know what I mean. So that's one kind of successful failure I'm pretty sure about.

Here's the second idea. Is it the actual failure that is the end in itself? Or is it the willingness to fail, the expectation of failure, the acceptance of failure, the desirability of failure? Can you imagine making failure desirable? Can you imagine aiming at failure? Can you appreciate making failure your goal?

You *can* if you have the right idea about the word failure—what I hope to convince you is the scientific version of failure. It is more than a stupid error, more than a shortcoming on your part, more than a miscalculation, more even than a chance to improve. Yes, more even than failures as life lessons. I know we all believe that a failure can be valuable if you learn something from it. After all, that's what we call experience. But how about a failure that does not aim at later self-improvement? How about failures that really are *an end in themselves*?

In this sense virtually all of science is a failure that is an end in itself. This is because scientific discoveries and facts are provisional. Science is constantly being revised. It may be successful for a time; it may remain successful even after it has been shown to be wrong in some essential way. That may seem strange, but good science is rarely completely wrong, just as it is never really completely right. The process is iterative. We scientists hop from failure to failure, happy with the interim results because they work so well and often are pretty close to the real thing.

Newton was famously wrong about two little things—time and space. They are not absolute. Gravity is not explained by the attraction between the centers of massive bodies, although it looks that way and can be usefully described that way. To the extent that we can explain it at all, it seems to be best understood, for now, as an emergent phenomenon of mass creating curvature in space. An imperfect but useful analogy is the way a heavy bowling ball on a mattress causes a depression and things placed on the mattress tend to fall toward it, as if they were being attracted to it. But Newton's failure in that one regard, even though it seems like a fundamental part of the theory of gravity, is not at all fatal to the success of his work. His equations quite accurately describe action at a distance between two bodies—sufficiently well to calculate how to dock a rocket with a space station orbiting some 250 miles away and moving at a speed of 17,000 miles per hour.

Nonetheless, there was a nagging inconsistency in Newton's model over what appeared as two different kinds of gravity. This inconsistency was what needled Einstein so much that he was ready to take a most unintuitive, illogical perspective. Although it's not exactly how Einstein thought about it, these two kinds of gravity are most easily experienced as the loss of gravity—weightlessness. One of them can be felt as distance from a massive body (the weightlessness experienced in outer space), and the other is due to acceleration (the weightless feeling you would have in a rapidly dropping elevator). They seem to be from two different and unrelated causes—the mass of a nearby body and the force resisted by inertia, or acceleration. Two hundred and fifty years later Einstein essentially corrected the failure of that part of Newtonian mechanics by showing that in the correct inertial frame, one that does not assume absolute time or space, the two kinds of gravity are the same.

Granted, it turned out to be a rather major correction, requiring a Copernican-sized shift in our point of view. But as with Copernicus it didn't require throwing everything else out. We continue to live our everyday lives in a Newtonian world where space and time seem sufficiently absolute, just as we continue to live most of our lives in a pre-Copernican world where the sun "rises" and "sets." That oversimplifies the story a great deal (see Notes), but the point is that Newton was successfully wrong and it was the very failed part of his model that led to Einstein's remarkable insights. Pretty good work.

A failure can be even less successful—that is, wholly incorrect—and still useful. An example from biology might be the longstanding principle known as "ontogeny recapitulates phylogeny." This tongue twister of a phrase, coined in 1866 by the "father of embryology," Ernst Haeckel, is simply a slightly bizarre attempt at making a complicated concept memorable by forming a jingle about it. It means that over the course of its development an embryo in the egg (or uterus) appears to go through all the stages of evolution of that organism. For example, mammals early in embryonic development have what appear to be gill-like structures, making them look a bit like fish. These structures eventually develop into our jaws and other muscle and bone groups of our heads and throats but have nothing directly to do with respiration, as gills do for fish. In fact, Haeckel's concept is completely wrong, even though it held sway for decades and led to many advances in embryology. Not only is it wrong about embryology, it is wrong about evolution. We didn't evolve from fishes (or apes for that matter); we shared a common ancestor that evolved into both of us, in the case of fish some 500 million years ago, and in the case of apes only about 85 to 90 million years ago.

Nonetheless, this failed ontogeny-phylogeny concept gave rise to important ideas about how development proceeds in clearly established stages, and that structures do evolve from earlier forms, possessing a common ancestry even if a contemporary divergence. Haeckel's work was painstaking and actually started the branch of science we today call embryology. In particular, he introduced comparative anatomy and development—that is, the notion that we can learn a great deal by making comparisons across species. This showed crucially that not only were species related but that their development proceeded in a similar way along certain principles. The value of this "failure" to modern biology cannot be overestimated. On the other hand, it remains damaging in that there are many people who still believe in it because they were taught it as schoolchildren. You remember, the silly business about having had a tail when you were an embryo.

You could object that Newtown's and Haeckel's failures eventually led to successes and were not therefore really ends in themselves. I think that's too much to ask of failure. Failures like these not only lead to greater insights, they often lead to very unpredictable insights. They force us to look at a problem differently because of the particular way in which they failed. This could be considered the case with Einstein's recognition that Newton's little failure was actually a fundamental misconception about time and space. We expect success to lead us to even greater success. What may not be so obvious is that failure can do the same.

These then are what I would call the failures that need no excuse, that stand shoulder to shoulder with success. They are the packing material, the innards, of science, and not giving them their full due is to miss more than half of what science is about and how it works. The big job I hope to do here is to remedy that.

. . .

There are many trivial things that can, and have, been said about failure. They are the kinds of aphorisms commonly found in Chinese restaurant fortune cookies. I'll sum them up in a paragraph and then we can get on to the interesting parts of it—the much wider and deeper functions of failure that are undeservedly ignored or, worse, thoughtlessly rejected as undesirable.

So then here we go: failing is part of succeeding. Failure builds character. Those who haven't failed haven't tried. You never know yourself until you've had a failure. You have to learn how to pick yourself up and get back in the game. And so on. I'm sure you can think of other, similar platitudes. And they're all okay advice, especially when you have someone on the phone who is really distraught over a recent failure in love or work or sport. Sure, failing is part of life and managing it is important for your happiness. And there are innumerable books loaded with mostly trivial advice about how to do all that. So let's us be done with it.

What we are interested in here, subtly but importantly different from those earlier instances, is where and when failure is actually an integral part of the process. Where it deserves to be right beside success, where it doesn't just make for an uplifting story of the young lad or lass who succeeds with perseverance, but where failure really has to be there for the process to occur properly. It is the difference between Edison (Thomas) failures and Einstein (Albert) failures. Edison claimed he never failed, just found 10,000 ways that didn't work. But eventually he succeeded. And of course it probably wasn't 10,000 wrong tries, but the actual number doesn't matter—it was a lot of trying and finally succeeding. This is good advice for an inventor, less so for a scientist. Einstein lived on failure, his own and those of others, not just ways that didn't work. His working failures were deep inconsistencies, failures of theory, failures that produced understanding even more than success. No failure, no science.

Now this is not true of all other grand human endeavors. You don't have to fail first in business to later become rich, you don't have to fail at writing to gain success as a novelist, and you don't have kill a few people to become a good doctor. In none of these endeavors is failure required—although it may happen, and unfortunately often does. Successful people may try to convince you that failing was a key to their success; they will espouse uplifting narratives to document it, and even write whole self-help books to aid you on your journey through failure. But it seems that way to them only retrospectively, because they failed and then they succeeded. The people who just succeeded right away don't have the kind of narrative that makes for good reading, and they rarely have any counsel you can use. It has been reported that when James Michener ("Tales of the South Pacific," 1947) was asked how to become a successful author, he replied, "Try to arrange to have your first novel turned into a musical by Rodgers and Hammerstein." Good advice.

Failure in all these endeavors is not uncommon, but it is not necessary. Not so in science. Failures are as informative as successes, sometimes more so, and, of course, sometimes less so. Failures may be disappointing at first, but successes that lead nowhere new are short-lived pleasures. *Conclusion* has a curious double meaning in science. We use it often as the heading of a section in our publications. We have the Methods and Results and, of course, the Conclusions (although "Conclusions" is now often called "Discussion," which seems humbler). In this context the word refers to what you can deduce or infer from the data—that is, what we have succeeded in finding out. But it also means *ending*, and we almost never mean, or want, that. Most of the time the "Conclusions" are themselves riddled with new questions. And many of those questions arise because some of the experiments didn't yield the expected results. They failed. Enrico Fermi, the pioneering nuclear physicist, would tell his students, "If your experiments succeed in proving the hypothesis, you have made a measurement; if they fail to prove the hypothesis, you have made a discovery."

In science you not only have to have the stomach for failure, you actually have to enjoy the taste of it.

If you accept the contention that failure is both an inevitable and a desirable part of science, then it is sensible to ask how much failure. After all, it can't be only failure, or at least I'm pretty sure it can't be only failure. But I think we commonly underestimate the amount of failure that is acceptable. Just to get some sense of the scale of failure let's look at it in other places and see what the ranges of tolerance are. We could begin with the natural world.

Nature's greatest predators—the kings of the jungle, the sea, the air; the killing machines of National Geographic specials—well, it turns out that they are successful on just about 7% of their attempted pursuits. You may think that a lion, killer whale, or red-tailed hawk can go out and bag some poor defenseless animal any time it gets a craving for a snack. In fact, 93% of the time they fail to capture their prey, which is why they have to be wily and nearly always on the hunt. Not only that, but they generally hunt around the edge of the herd, picking off the sick, the lame, and the old. There's a reason for that: the failure rate for bagging some nice young and still juicy creature is even higher. Nonetheless, we still think of them as being at the top of the food chain and anoint them as kings of their niches. It seems from the biological perspective you can put up with a lot of failure and still make a decent living. (I suppose you could turn this around and say that the *prey* are successful a remarkable 93% of the time, but that is a difficult calculus since a prey animal can fail only once. And scientists are hunters, not prey, I hope.)

Evolution itself is a marvel of failure. Well over 99% of the species that have ever made an appearance are now extinct. Species continue to go extinct, some scientists believe at a currently alarming rate. How, out of all this failure, could the marvelously complex creatures that we observe have emerged? Could all of these remarkable animals and plants and ecosystems have been created by failure? Perhaps it is not so hard to see the attraction of creation narratives and why they elicit widespread belief when the alternative explanation is failure. But, like it or not, that is how it works.

Darwin's great insight was that evolution proceeds by random changes in an organism's makeup, followed by a selection process that favors beneficial changes. It thereby wipes out over time the useless or harmful changes and even the status quo. Today we understand that these changes are due to mutations in genes and these mutations are overwhelmingly failures. The failures perish, most of them immediately, the less failed ones after hundreds of thousands or even millions of years. But eventually they fail. A billion or more years of evolution is primarily a record of failures.

And it doesn't end there. The actual mechanism for evolution—the nuts and bolts of it, what we glibly call random mutation—depends itself on failure. Sperm and egg cells copy DNA from the parents to the offspring. DNA is a molecule with a structure famously identified as a double helix, like two spiral staircases wrapped around each other. It's the double part that is the key to DNA's hereditary function. Each of the two helices is a copy of the other. If they split apart, as they do in egg and sperm cells, then each helix, using some enzymes and other chemicals inside the cell, can fashion a new partner helix for itself. That is, it can copy itself. But the copying process is not flawless; it makes mistakes. These mistakes are what we call random mutations. They are random because the chemical process of copying is simply imperfect; it doesn't favor any particular kind of mistake. Some of the mistakes result in changes to a gene, one section of the DNA molecule, that are fortuitously beneficial, and the resulting offspring with this improved gene has some advantage over others with the older model gene. Stronger, faster, smarter, whatever. But since the process is essentially a copying mistake, a failure, most of the time the change is harmful, or at best useless. Without this faulty copy mechanism there would be no evolution, nothing for natural selection to work on. From what can be seen only as an overwhelming tide of mistakes and failures, the living world emerged. All the complexity, all the apparently clockwork precision of life, from the developing embryo to the most elaborate ecosystem—all of it is due to failure at an almost unimaginable scale. When you have a few billion years to mess about you can put up with a lot of failure.

Perhaps a more pedestrian, but more immediate, illustration might be athletics. This is an activity where success seems to be important, and failure is to be avoided. So, how acceptable is failure in sports? In the game of baseball a position player's salary—as opposed to a pitcher's—is generally keyed to their batting average. This number is the percentage of times the player reaches base safely by hitting the ball. It is calculated by dividing the number of hits by the number of at-bats (the number of times the player came to bat and therefore the number of opportunities to make a hit). Because the baseball season goes on for so long and player careers often last a dozen or more seasons, this batting average can be calculated to three significant decimal places. Thus the famous Joe DiMaggio of the Yankees had a lifetime average of 0.325. The decimal point is typically dropped, and one says DiMaggio had a "325 lifetime batting average." Joe DiMaggio was one of the all-time greatest baseball players, and perhaps along with Ted Williams of the rival Boston Red Sox (344 lifetime average), they were the greatest hitters ever. But what their averages tell us is that nearly 7 out of 10 times that they came to bat . . . they failed. They struck out, grounded out, flied out, or were out in any of a number of ways, and they simply went back to the dugout and sat down until their next chance to bat.

(To be precise, there was also the possibility of a walk, which allows the batter to advance to a free base because the pitcher threw the ball out of the strike zone four times and the batter didn't swing. But in baseball statistics a walk doesn't count as a batting opportunity and so does not affect the batting average. In fact, Ted Williams walked a remarkable 2,021 times, while DiMaggio walked only 790 times. There are many complicated reasons behind these numbers, some of which have to do with player skills and some of which are more a matter of strategy and other subtle reasons that are not so relevant to this argument. Lest you worry that I am about to slip into a long polemic about some baseball minutiae, I will resist and stop here.)

DiMaggio had 6,821 plate appearances during his 13 seasons, and he hit safely in 2,214 of those. But he was out 4,607 times. Nearly twice the number of times he was safe. Even more impressively, Ted Williams, in 19 seasons, came up to bat 7,706 times, hit safely 2,654 times, but was out 5,052 times. The two greatest hitters in the history of baseball had a combined total of nearly 10,000 failures!

Only a few players hit regularly at 300 and above, and these are the highest paid players in the game, commanding salaries in excess of $10 million per year. Ten million a year to fail 7 out of 10 times, dependably. Clearly failure can be a good bet.

So what's the answer? How much failure is acceptable? Of course there's no number or precise quantity to be calculated. But we can see from just these few examples that the acceptable rate is likely much higher than you would have imagined. In the limit, success need occur only once; failures can occur again and again as long as your resources, or your life, doesn't run out. And even a dependable failure rate of 80–90% could be considered successful. Experience, that much-valued attribute of the learned, is after all the result of not getting it right the first time. Niels Bohr described an expert as "a person who has made all the mistakes that can be made in a very narrow field." Notice that it is not someone who has been a success in some narrow field.